Welcome To

ISSCC 2014
Session 6
Technologies for High-Speed Data Networks

Memory and System Architecture for 400Gb/s Networking and Beyond

Dinesh Maheshwari

CTO, Memory Products Division

Cypress Semiconductor

ISSCC 2014

OUTLINE

INTRODUCTION

MEMORY METRIC – NETWORKING vs. COMPUTE

NETWORKING MEMORY APPLICATIONS & REQM'TS

PROBLEM FOR HIGH PERFORMANCE NETWORKING

SOLUTION - 2.5D SRAM INTEGRATION

ISSCC 2014 Theme
SILICON SYSTEMS BRIDGING THE CLOUD

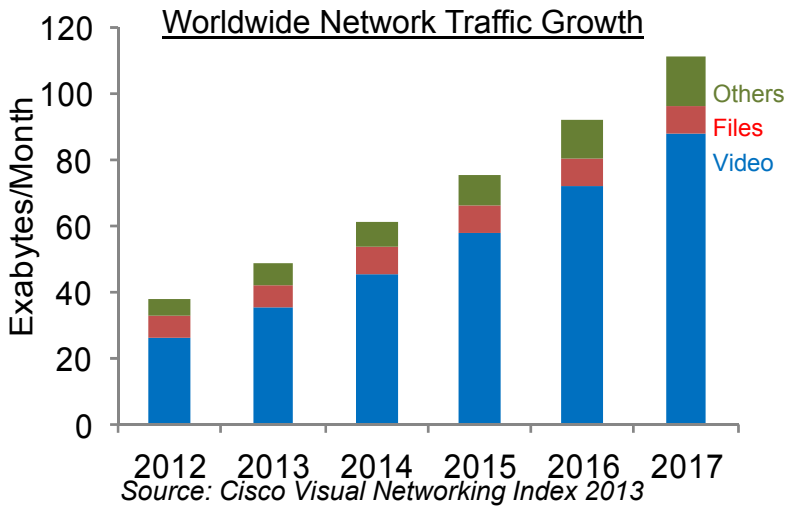

Source: Cisco Visual Networking Index 2013

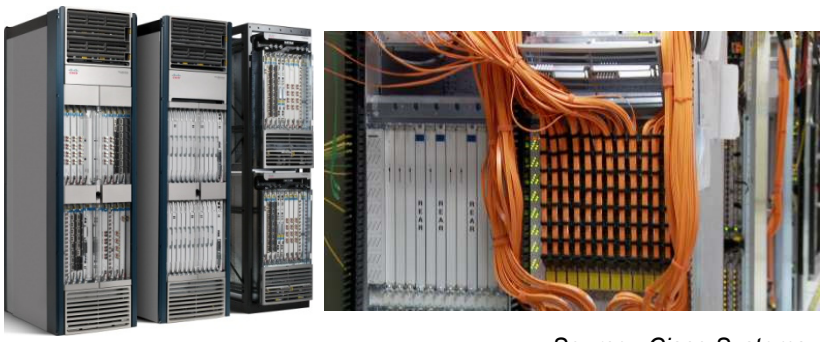
Source : Cisco Systems

Networks Build and Bridge the Clouds

Switches & Routers are at the heart of Networks

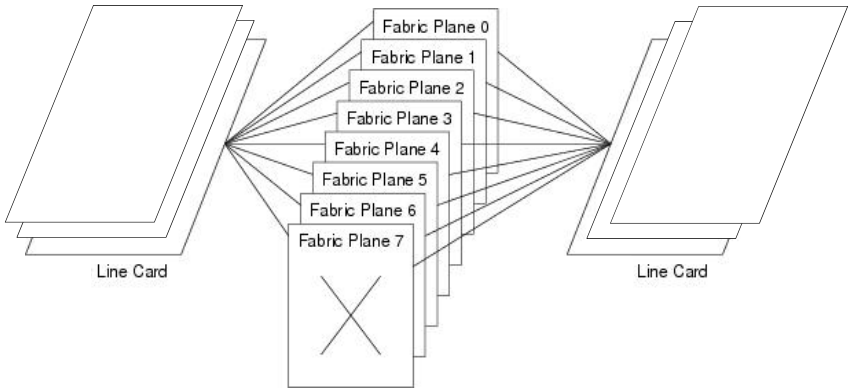

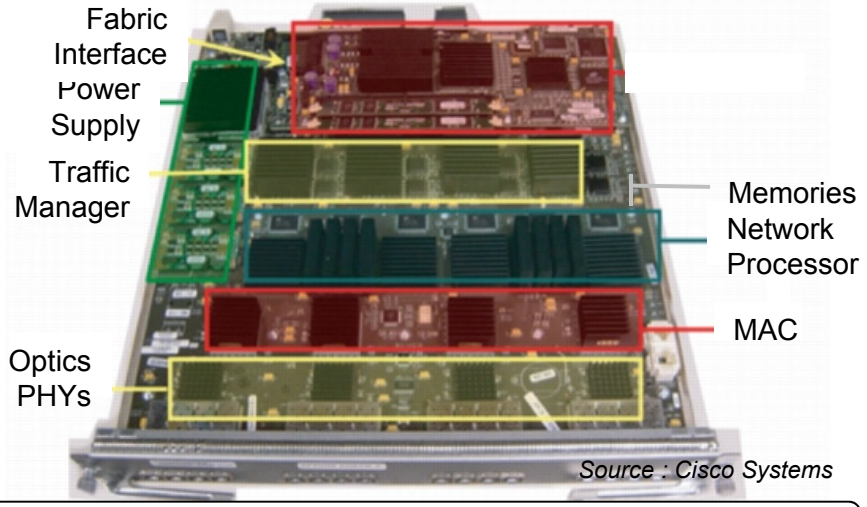

Switches = Line Cards + Switch Fabric

A Switch Line Card

NETWORKING LINECARD vs SERVER MOTHERBOARD

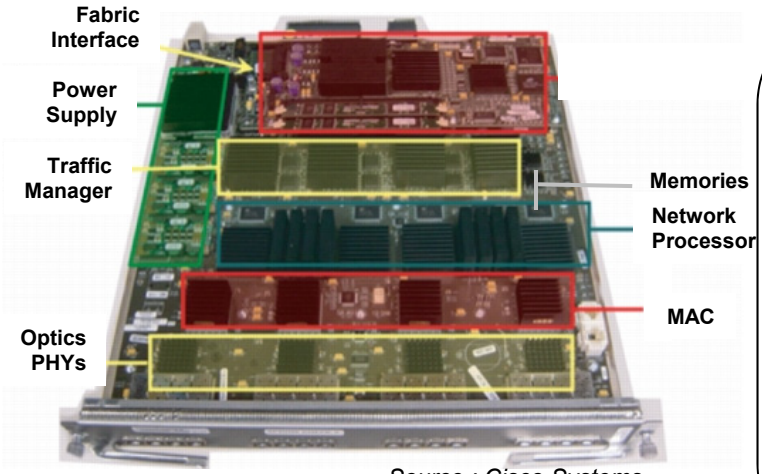

Source : Cisco Systems

Networking Linecard Metric

Packets Per Second

Corresponds to # of min size packet in switched BW (e.g. 100gbps)

Unit - MPPS

Discrete values -150 MPPS, 240 MPPS, 300 MPPS etc.

149 MPPS line card does not exist

Memory Access Pattern - Random

Memory Metric - Random Transaction Rate (MT/S)

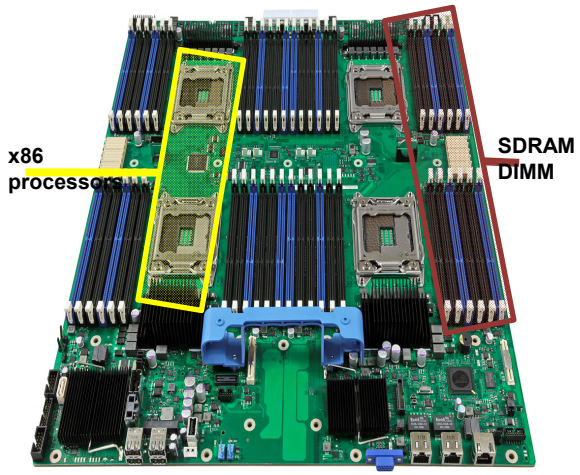

Source : Intel

Compute Server Metric

Instructions Per Second

Unit - MIPS

Continuous values

999 MIPS can coexist with 1000 MIPS

Memory Access Pattern – Spatial & Temporal Locality

Cache Memory Metric – Latency (ns)

Main Memory Metric – BW (Gb/s), Density (Gb)

NETWORKING LINECARD - MEMORY APPLICATIONS

SIX KEY DATA PLANE MEMORY FUNCTIONS

FORWARDING LOOKUP
CLASSIFICATION LOOKUP
STATISTICS UPDATE
STATE UPDATE
PACKET BUFFER
PACKET/CELL SCHEDULING

MAC/Framer
- Physical layer MAC
- (De)Frame data
- Over Subscription (OS) Packet Buffer
- Queue w customer policy

Network Processor (NPU)- ASIC/ASSP
- Parse & Identify flow
- Classify
- Forward
- Update states/statistics
- Mark QoS/Add Header

Ingress Traffic Manager (TM)
- Police
- Manage congestion
- Buffer Cell
- Queue Cell

Switch Fabric (SF)
- Switch cells from Input to Output

Egress Traffic Manager (TM)
- Reassemble packets
- Shape traffic
- Buffer egress traffic
- Queue egress traffic
- Multi-cast buffering

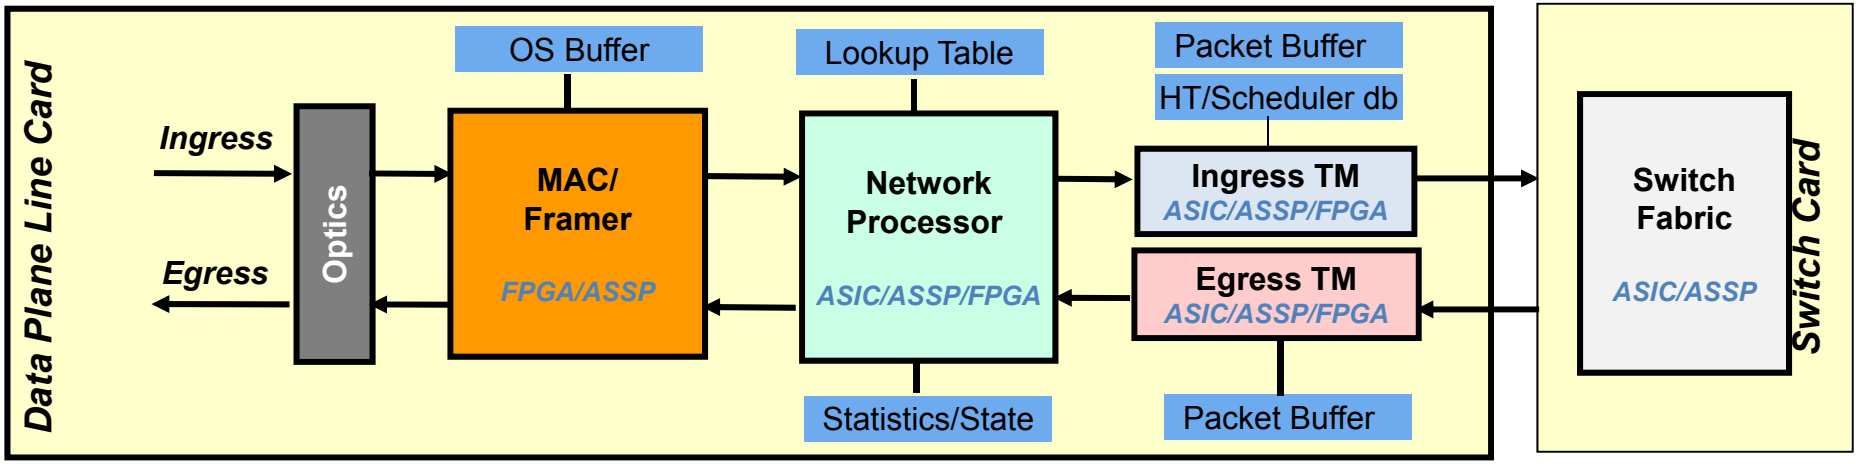

BEYOND 2014: NETWORK TRENDS

LINE RATE GROWTH ACCELERATES

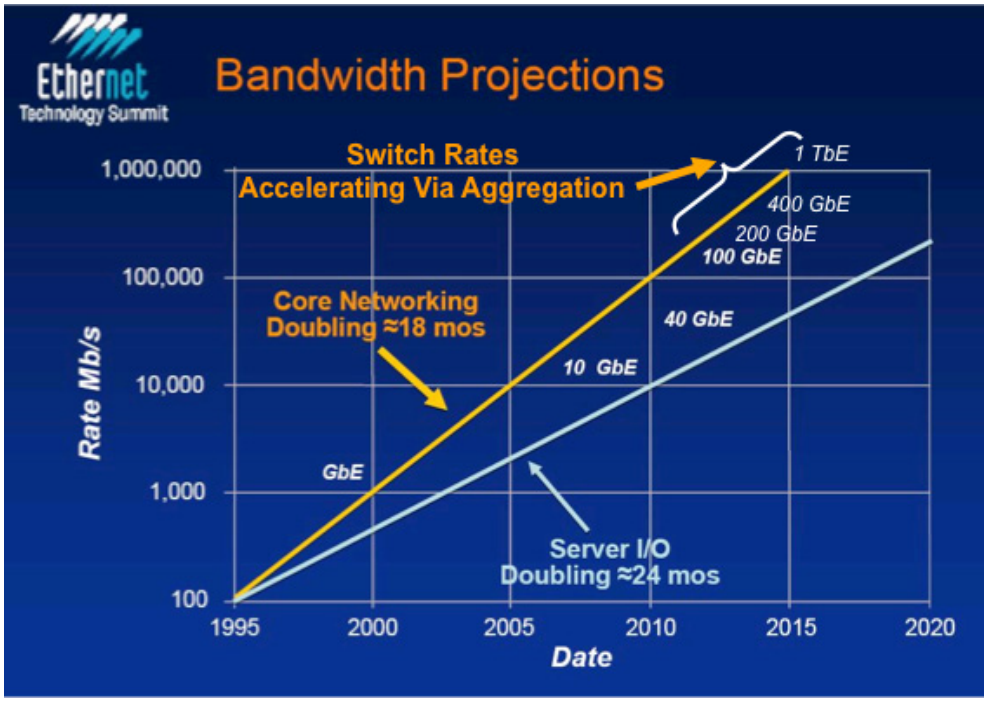

Source : Ethernet Technology Summit

IMPLICATIONS FOR MEMORY

Network Metric	=	Packet Rate - R (MP/S)
Memory Metric	=	Random Transaction Rate (RTR) Rate Of Truly Random Memory Accesses (MT/S)
RTR Required	=	Packet Rate (R) X Accesses/Packet
Network Memory Wall	=	Packet Rate Is Limited By Memory RTR

NETWORKING MEMORY SYSTEM REQUIREMENTS

WAN/MAN Networking memory Requirements						
RTT = Round Trip Time = 250ms Lf = Latency Factor = 3 R= Packet Rate (MP/s) Q=TM queues r = # of routes F = # of Flows		Formulaic Req'mts		100G	200G	400G
		Requirements	Atomic Requirement	R=150 Q=1.25K r=0.5M F=0.5M	R=300 Q=2.5K r=1M F=1M	R=600 Q=5K r=2M F=2M
PACKET BUFFER (SDRAM)	RTR (MT/s)	n/a	(1r+1w)R			
	B/W (Gb/s)	2R*84B	(1r+1w)R*84B	200	400	800
	DENSITY (Gb)	R*84B*RTT		25	50	100
HEAD-TAIL /SCHEDULER (SRAM)	RTR (MT/s)	4R	(1r+1w)2R	600	1200	2400
	B/W (Gb/s)	4R*84B	(1r+1w)2R*84B	400	800	1600
	DENSITY (Mb)	Q*2R*64B*tRC*Lf		37.5	150	600
LOOKUP (SRAM)	RTR (MT/s)	4R-8R	(1r/w)R	600-1200	1200-2400	2400-4800
	B/W (Gb/s)	RTR*64	(1r/w)R*64	38.4-76.8	76.8-153.6	153.6-307.2
	DENSITY (Mb)	~(160Mb/M entries)*r		80	160	320
STATISTICS +STATES (SRAM)	RTR (MT/s)	2R*(2-8) Updates	(1r+1w)R	600-2400	1200-4800	2400-9600
	B/W (Gb/s)	RTR*64	(1r+1w)R*64	38.4-153.6	76.8-307.2	153.6-614.4
	DENSITY (Mb)	64*C*F, C=4		128	256	512

Networking memory RTR, BW and density doubles with every generation

CORE RTR FOR DIFFERENT MEMORIES

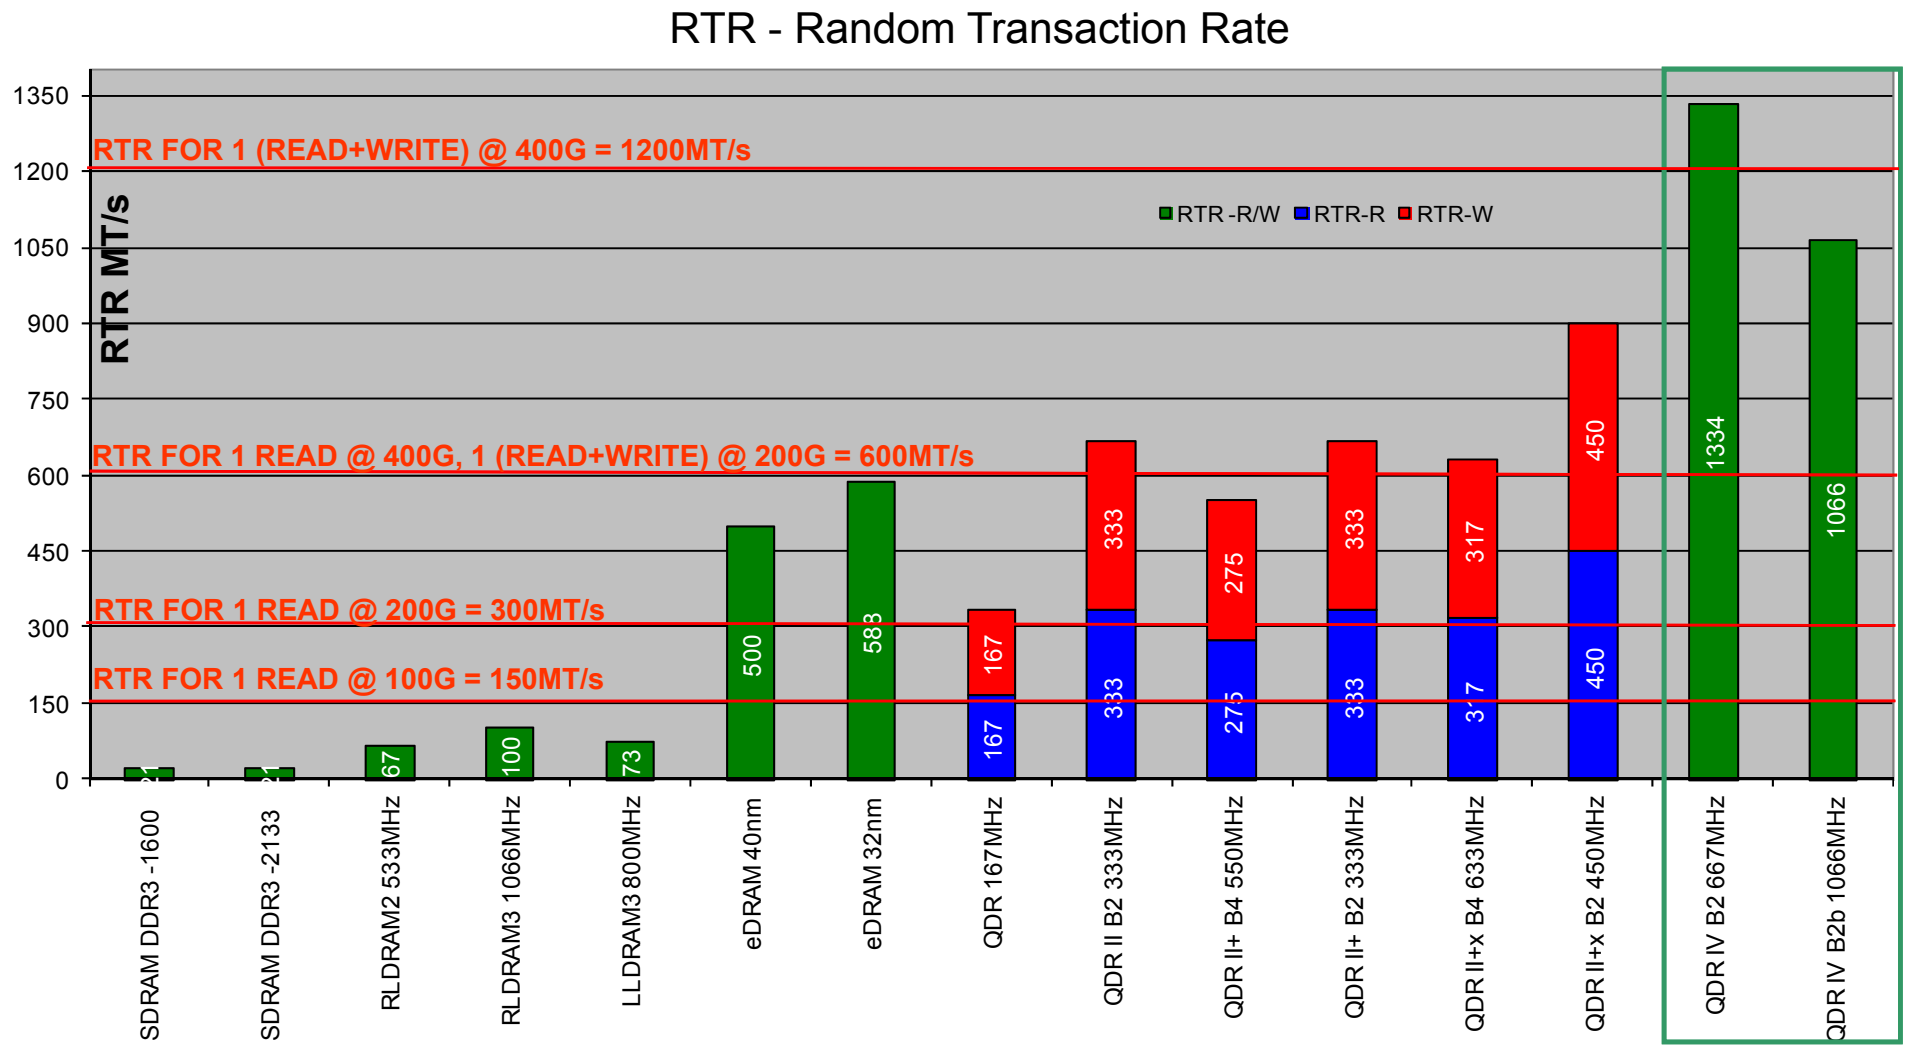

Only SRAMs can scale to serve the requisite RTR

NEED 2x IMPROVEMENT IN READ & WRITE ATOMIC OPERATIONS EVERY GENERATION

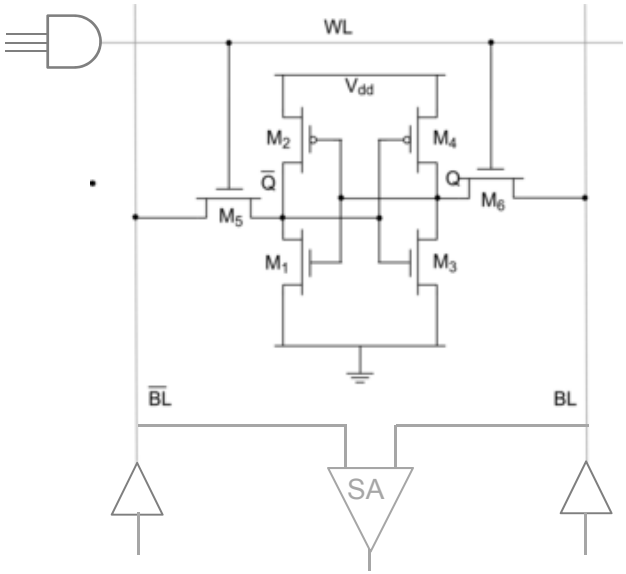

Core RTR of memory is limited by atomic operations with largest cycle time, everything else can be pipelined

Limiting atomic operations are

1. Wordline activation to sense amp out during read
2. Wordline & driver activation to cell write during write

RAM cell & sense amp architecture, read & write assist circuits directly impact RTR and thereby the linecard packet rates

Need to target 2x improvement in read & write atomic operations from one generation to next

CURRENT SRAM SOLUTIONS FOR <= 200G NETWORKING

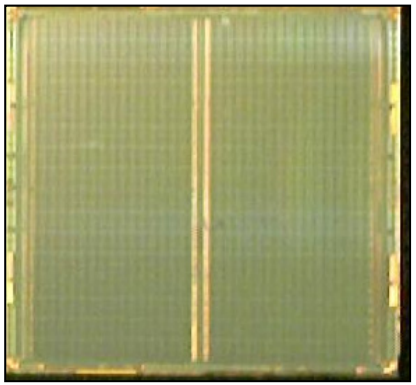

QDR-II+

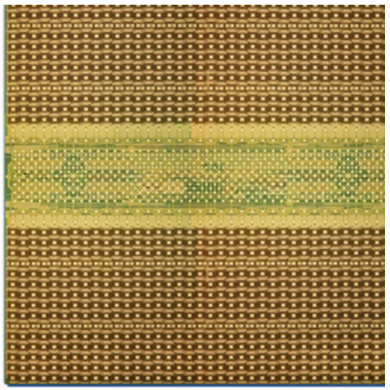

QDR-IV

Key Information	
Technology	65nm
Core / IO VDD	1.2V / 1.2V
SRAM Cell	0.525um^2

Products	Process Tech	RTR (MT/s)	Optimization	Status
QDR-II+	65nm	666		Shipping
QDR-IV	65nm	2132	Memory array supporting high RTR : Micro-architecture change	Sampling

QDR-IV can support 3.2x higher packet rate than QDR-II+ - at the same process node

SDRAM IS EXPENSIVE FOR LOOKUP

BASICS: DRAM LIMITATIONS

tRC: Random Access Cycle time:

Min time between random read or write accesses

→ Max of one random access every tRC

tFAW: Four Access Window Time

Min time in which four activates are allowed

→ Max of four banks can be open within tFAW

Interface width

Max # of bits per each access

Current Product Example: DDR3-1600

tRC= 48.75ns; tFAW=30ns, interface width = x16

SDRAM IS EXPENSIVE FOR LOOKUP
100G, 1 x32 READ/PACKET USING DDR3-1600

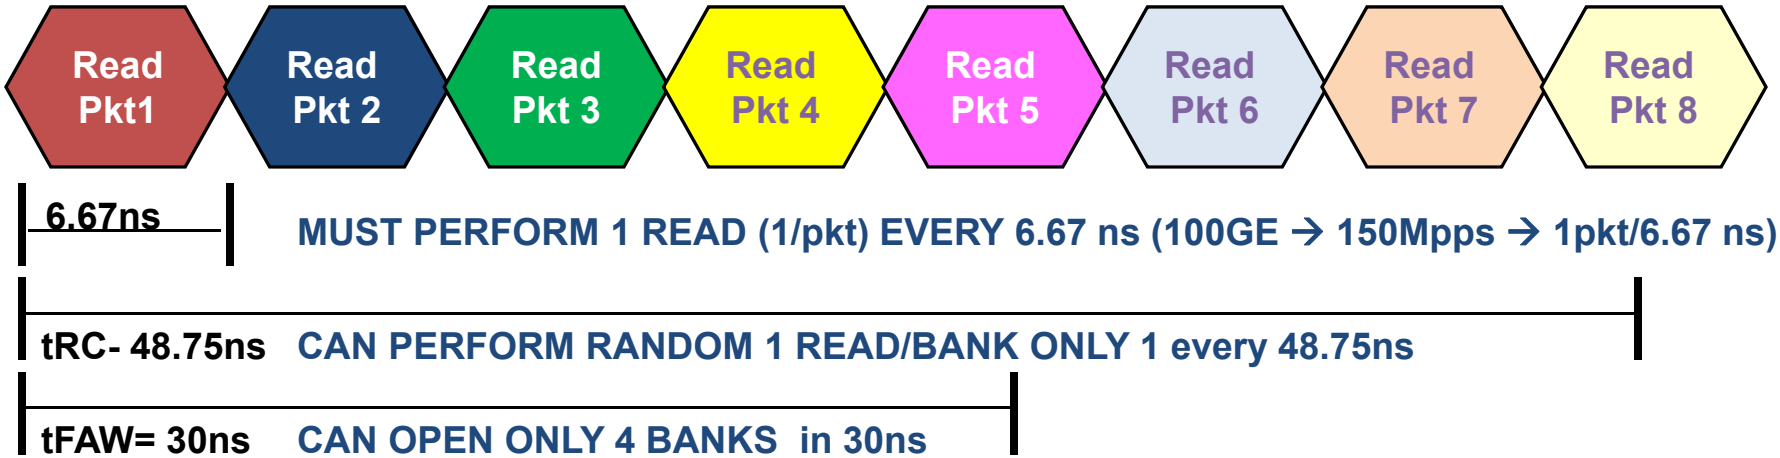

Problem 1: Need random read every 6.67ns;but can only do one every 48.75ns(tRC)
Solution: Replicate (multiple copies) data in banks need 8x replication (48.75/6.67)

Problem 2: Can open max four banks per chip simultaneously (TFAW).
Solution: Use two chips for 8x replication (but waste ½ banks in each chip)

Problem 3: Have to perform x32 read, but each chip is x16
Solution: Parse access into two x16 reads → Now 4 chips required

SDRAM IS EXPENSIVE FOR LOOKUP
100G, 1 x32 READ/PACKET USING DDR3-1600

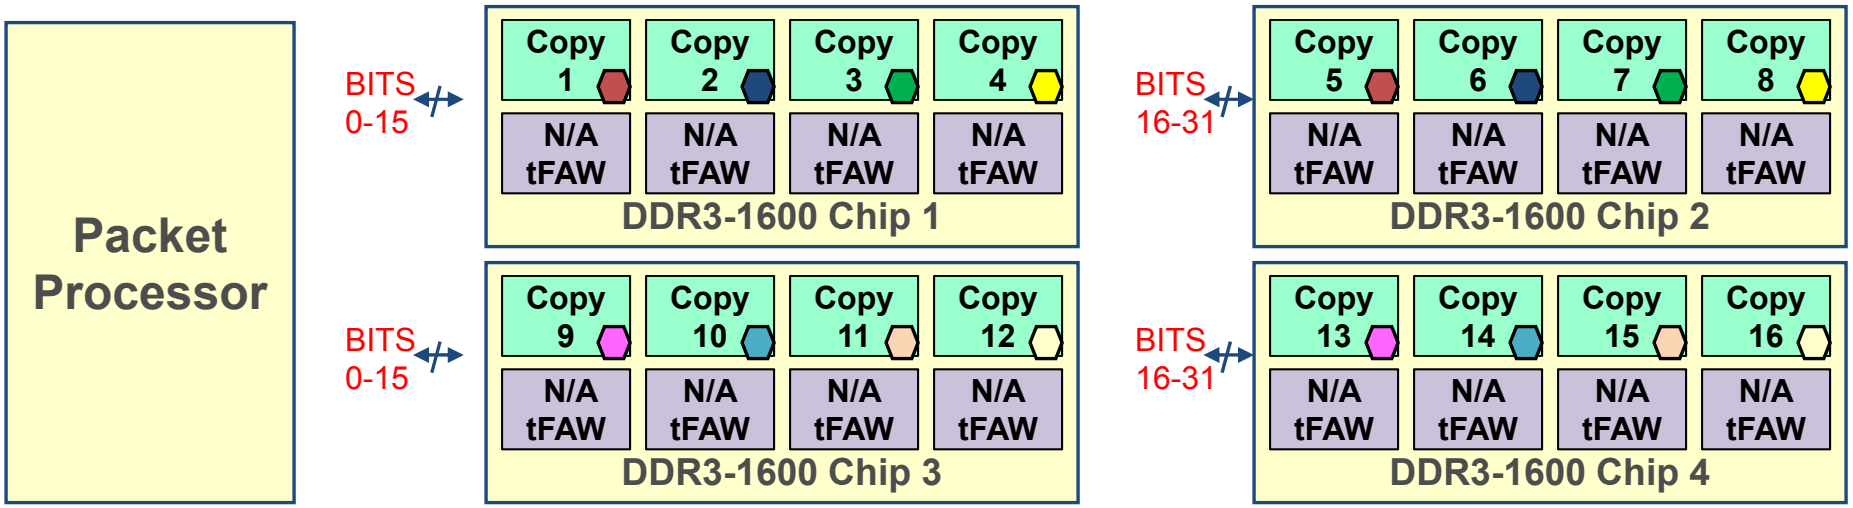

Four chips required for 1 read/packet @ 100G

Real world forwarding look up requires 4-8 reads/packet

100G look up requires 16-32 separate chips!!

Adds component cost, pincount, and power

SDRAM + SRAM CAN SERVE PACKET BUFFER

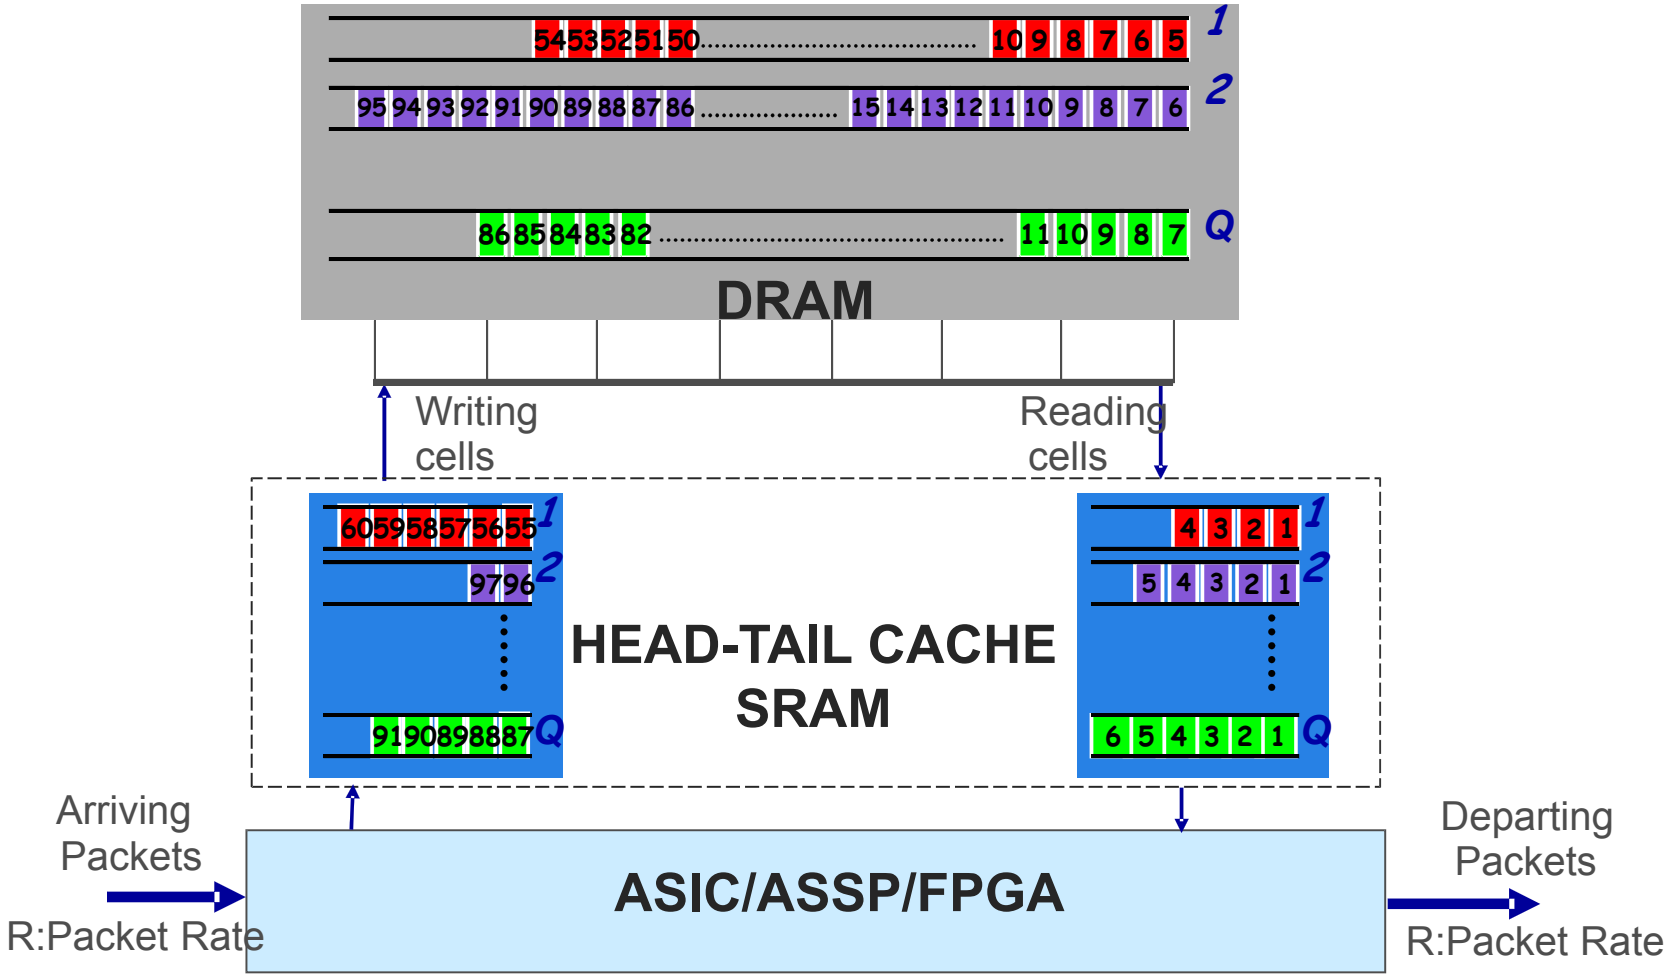

SRAM (HEAD-TAIL CACHE) RTR ENABLES THE USE OF SDRAM FOR DENSITY

SOLUTION FOR 100G, 200G - SRAM EXTERNAL OR ON DIE

WAN/MAN Networking memory Requirements						
RTT = Round Trip Time = 250ms Lf = Latency Factor = 3 R= Packet Rate (MP/s) Q=TM queues r = # of routes F = # of Flows		Formulaic Req'mts		100G	200G	400G
		Requirements	Atomic Requirement	R=150 Q=1.25K r=0.5M F=0.5M	R=300 Q=2.5K r=1M F=1M	R=600 Q=5K r=2M F=2M
PACKET BUFFER (SDRAM)	RTR (MT/s)	n/a	(1r+1w)R			
	B/W (Gb/s)	2R*84B	(1r+1w)R*84B	200	400	800
	DENSITY (Gb)	R*84B*RTT		25	50	100
HEAD-TAIL /SCHEDULER (SRAM)	RTR (MT/s)	4R	(1r+1w)2R	600	1200	2400
	B/W (Gb/s)	4R*84B	(1r+1w)2R*84B	400	800	1600
	DENSITY (Mb)	Q*2R*64B*tRC*Lf		37.5	150	600
LOOKUP (SRAM)	RTR (MT/s)	4R-8R	(1r/w)R	600-1200	1200-2400	2400-4800
	B/W (Gb/s)	RTR*64	(1r/w)R*64	38.4-76.8	76.8-153.6	153.6-307.2
	DENSITY (Mb)	~(160Mb/M entries)*r		80	160	320
STATISTICS +STATES (SRAM)	RTR (MT/s)	2R*(2-8) Update	(1r+1w)R	600-2400	1200-4800	2400-9600
	B/W (Gb/s)	RTR*64	(1r+1w)R*64	38.4-153.6	76.8-307.2	153.6-614.4
	DENSITY (Mb)	64*C*F, C=4		128	256	512

Packaged SRAM used in 100G and some 200G Line Cards

Some ASIC/ASSPs embedded SRAM on die for better footprint, SI, mW/Gbps & cost

~500mm^2 with up to ~60% comprised of SRAM

Lower cost segments @ same packet rate require new die with less SRAM

higher NRE for family

MEMORY SOLUTIONS FOR NEXT GEN 400G NETWORKING

R = 600MP/s	WAN/MAN (400G)								
	Effective RTR		Density	RLDRAM+eRAM		SDRAM+RLDRAM		SDRAM+SRAM	
	Formulaic	MT/s	Formulaic	MEM Gb	COST AU	MEM Gb	COST AU	MEM Gb	COST AU
PACKET BUFFER			R*84B*250ms	RLDRAM 100	100* (25-50)	SDRAM 100	100	SDRAM 100	100
HEAD-TAIL CACHE	(4R-6R)	2400-3600	3Q*2R*64B* tRC	SRAM 0.24	24	RLDRAM 1.2*12	1.2 *12 *(25-50)	SRAM 1.2	1.2*100
LOOKUP	(4R-8R)	2400-4800	160Mb/1M 1M-2M	SRAM 0.16-0.32	16-32	RLDRAM (.16-.32) *6	(.16-.32) *6 *(25-50)	SRAM 0.16-0.32	16-32
STATISTICS	2R*(2-8) Updates	2400-9600	64*#C*#F C=4 F=1M-2M	SRAM .256-.512	25.6-51.2	RLDRAM (PASSIVE) .256-.512	(0.256-0.512) *(25-50)	SRAM .256-.512	25.6-51.2
TOTAL EXT MEM COST					2565-5107		490-941		261-303

Q = 10K

SDRAM+SRAM SOLUTION BECOMES EVEN MORE ATTRACTIVE AS Q INCREASES

COST OF 1 Gb SDRAM	= 1 AU
COST OF 1Gb RLDRAM/LLDRAM	= 25-50AU
COST OF 1Gb SRAM	= 100AU

Cost effective, high performance memory subsystem for 400G line card -
95%-98% low cost SDRAM bits for packet buffer and
2%-5% SRAM for lookup, statistics, scheduling & head tail cache

SRAM DENSITY - 16nm vs 28nm

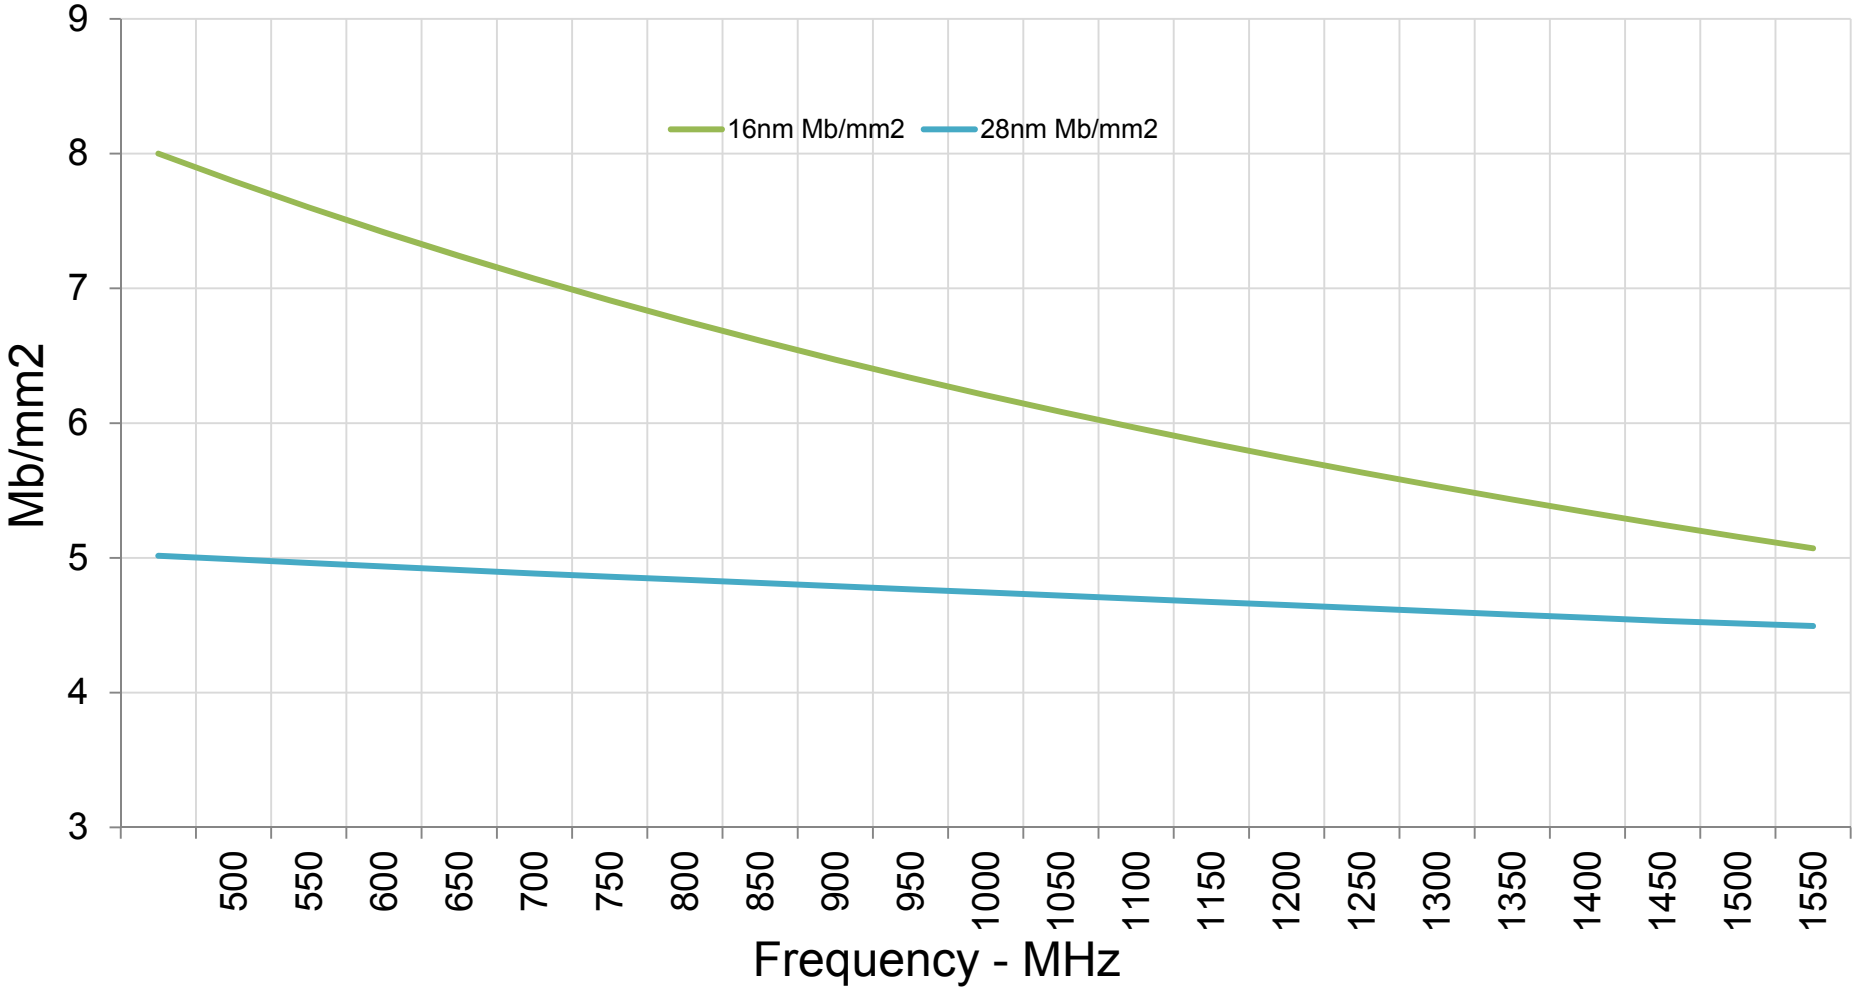

Memory density at 1500MHz and above scales by ~1.1x or less from 28nm to 16nm

ON DIE SRAM IS IMPRACTICAL FOR ≥400G

WAN/MAN Networking memory Requirements							Interface Power	
RTT=Round Trip Time=250ms Lf = Latency Factor = 3 R= Packet Rate (MP/s) Q=TM queues r = # of routes F = # of Flows		Formulaic Req'mts		100G	200G	400G	400G	
		Requirements	Atomic Requirement	R=150 Q=1.25K r=0.5M F=0.5M	R=300 Q=2.5K r=1M F=1M	R=600 Q=5K r=2M F=2M	Serial (W)	HBM (W)
PACKET BUFFER (SDRAM)	RTR (MT/s)	n/a	(1r+1w)R				4	0.8
	B/W (Gb/s)	2R*84B	(1r+1w)R*84B	200	400	800		
	DENSITY (Gb)	R*84B*RTT		25	50	100		
HEAD-TAIL /SCHEDULER (SRAM)	RTR (MT/s)	4R	(1r+1w)2R	600	1200	2400	8	1.6
	B/W (Gb/s)	4R*84B	(1r+1w)2R*84B	400	800	1600		
	DENSITY (Mb)	Q*2R*64B*tRC*Lf		37.5	150	600		
LOOKUP (SRAM)	RTR (MT/s)	4R-8R	(1r/w)R	600-1200	1200-2400	2400-4800	0.75-1.5	0.15-0.3
	B/W (Gb/s)	RTR*64	(1r/w)R*64	38.4-76.8	76.8-153.6	153.6-307.2		
	DENSITY (Mb)	~(160Mb/M entries)*r		80	160	320		
STATISTICS +STATES (SRAM)	RTR (MT/s)	2R*(2-8) Update	(1r+1w)R	600-2400	1200-4800	2400-9600	0.75-3	0.15-0.6
	B/W (Gb/s)	RTR*64	(1r+1w)R*64	38.4-153.6	76.8-307.2	153.6-614.4		
	DENSITY (Mb)	64*C*F, C=4		128	256	512		

For 400G ASSP/ASIC, need to double the SRAM density

- But density improvement from process node N to N+1 is not 2x anymore but by 1.1x
- 500mm^2 die in 28nm for 200G with 60% SRAM ported to 16nm for 400G will be ~ 745mm^2
- Die size close to reticle limit - exacerbates yield & cost of lower end segments

At > 400G, embedding all the SRAM would make the die size bigger than reticle limit

THE SOLUTION - SRAM INTEGRATED IN 2.5D PACKAGE

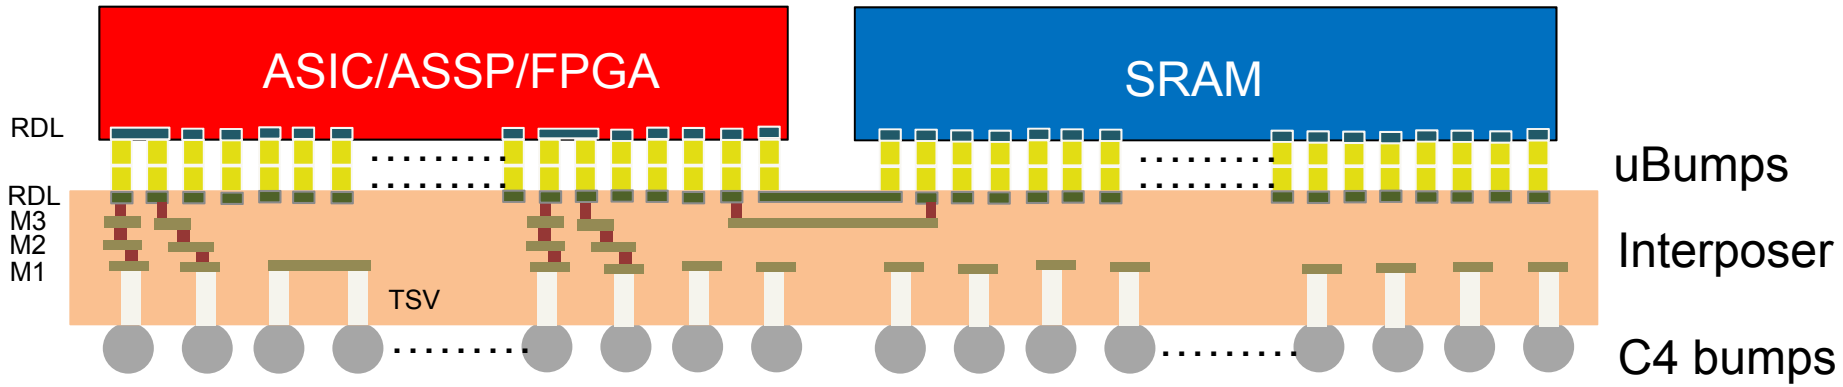

SOLUTION VALUE PROPOSITION

All the benefits of integration (small footprint, simple board design, low mW/Gbps) while…

Enabling cost effective family of parts for different segments at low NRE

Preserving TM/NPU die space for differentiating L3 to L7 processing functions

Improving TM/NPU die yield

SRAM

RTR: 6-12 GT/s

B/W: 1.6 Tbps

INTERFACE: HBM based
(Ddefined by JEDEC for HBM SDRAM)

I/O POWER: ~1 mW/Gbps

1-4 Slices in Package

SRAMs + ASIC/ASSP/FPGA IN 2.5D

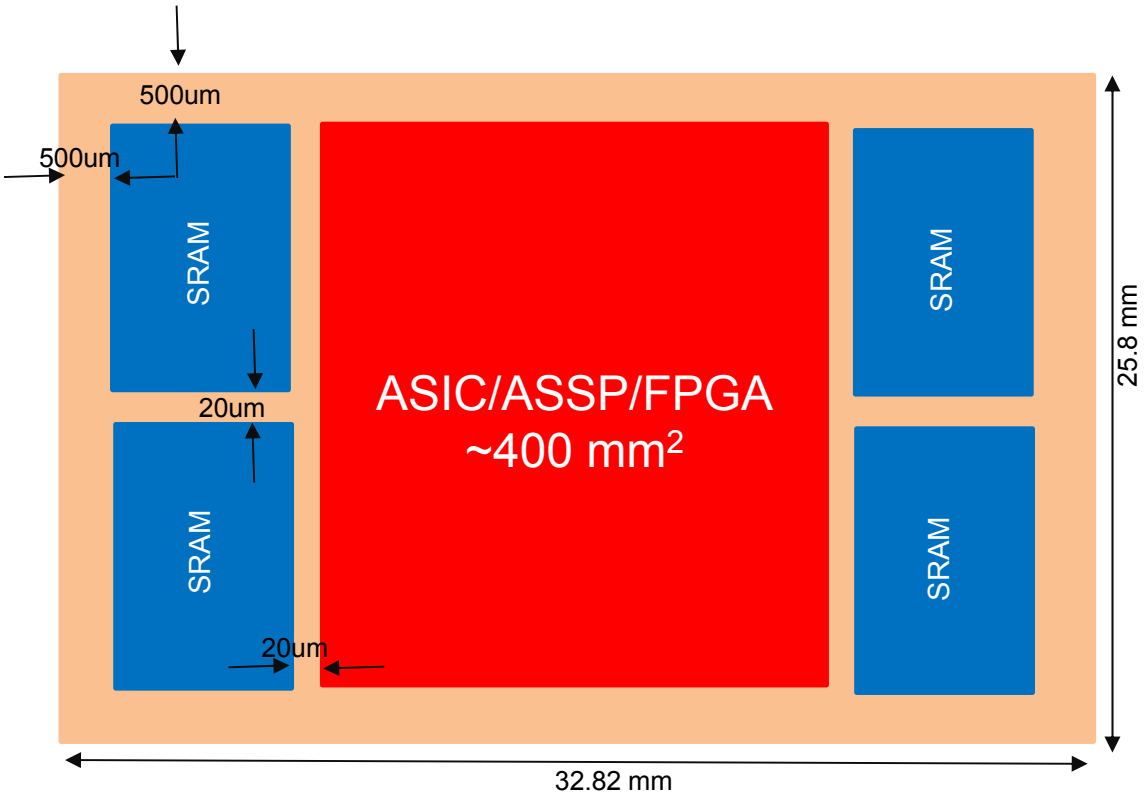

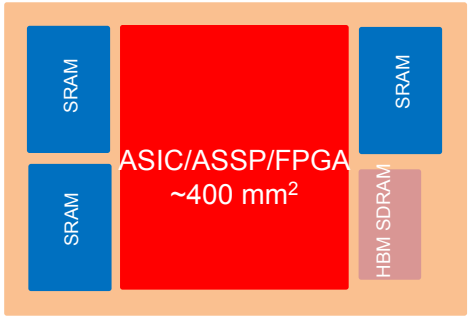

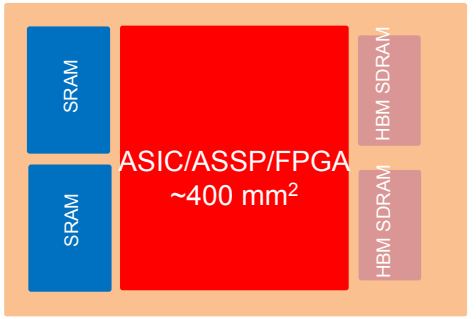

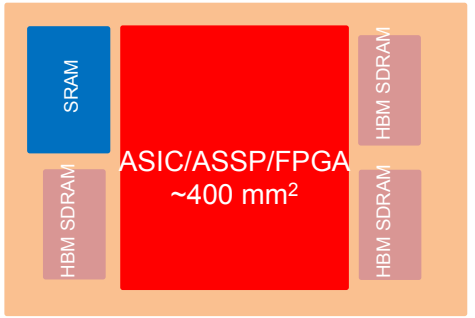

4 HBM Interfaces - "mix-n-match" SRAM & SDRAM for multiple SKUs @ lower NRE

HBM interface – 8 Channels of 128 I/Os DDR; Area - 3mm^2 Diffusion, 7mm^2 Top Metal

4 HBM interfaces on ASIC/ASSP/FPGA –12mm^2 total diffusion area

4 equivalent interface using SERDES would take 100mm^2

CONCLUSION

Line Card rates continue to accelerate to meet Internet traffic demands

SRAMs continue to be the best memory to deliver the RTR that doubles for each new generation of Line Cards

2.5D integration of SRAMs with ASIC/ASSP/FPGA will be the solution for next generation networking systems

High-Capacity Scalable Optical Communication for Future Optical Transport Network

Feb. 10th , 2014

NTT Network Innovation Laboratories

Yutaka Miyamoto and Masahito Tomizawa

OUTLINE

- **Background**
  - Today's Optical Transport Network
- **100Gbps / channel Digital Coherent Transport**
  - Proof of concept : **100G DSP ASIC**
- **Key Issues for Beyond-100G / channel Transport**
- **Future evolution beyond 1 Pb/s capacity by Space Division Multiplexing**
- **Conclusion**

High-Capacity Optical Transport Network

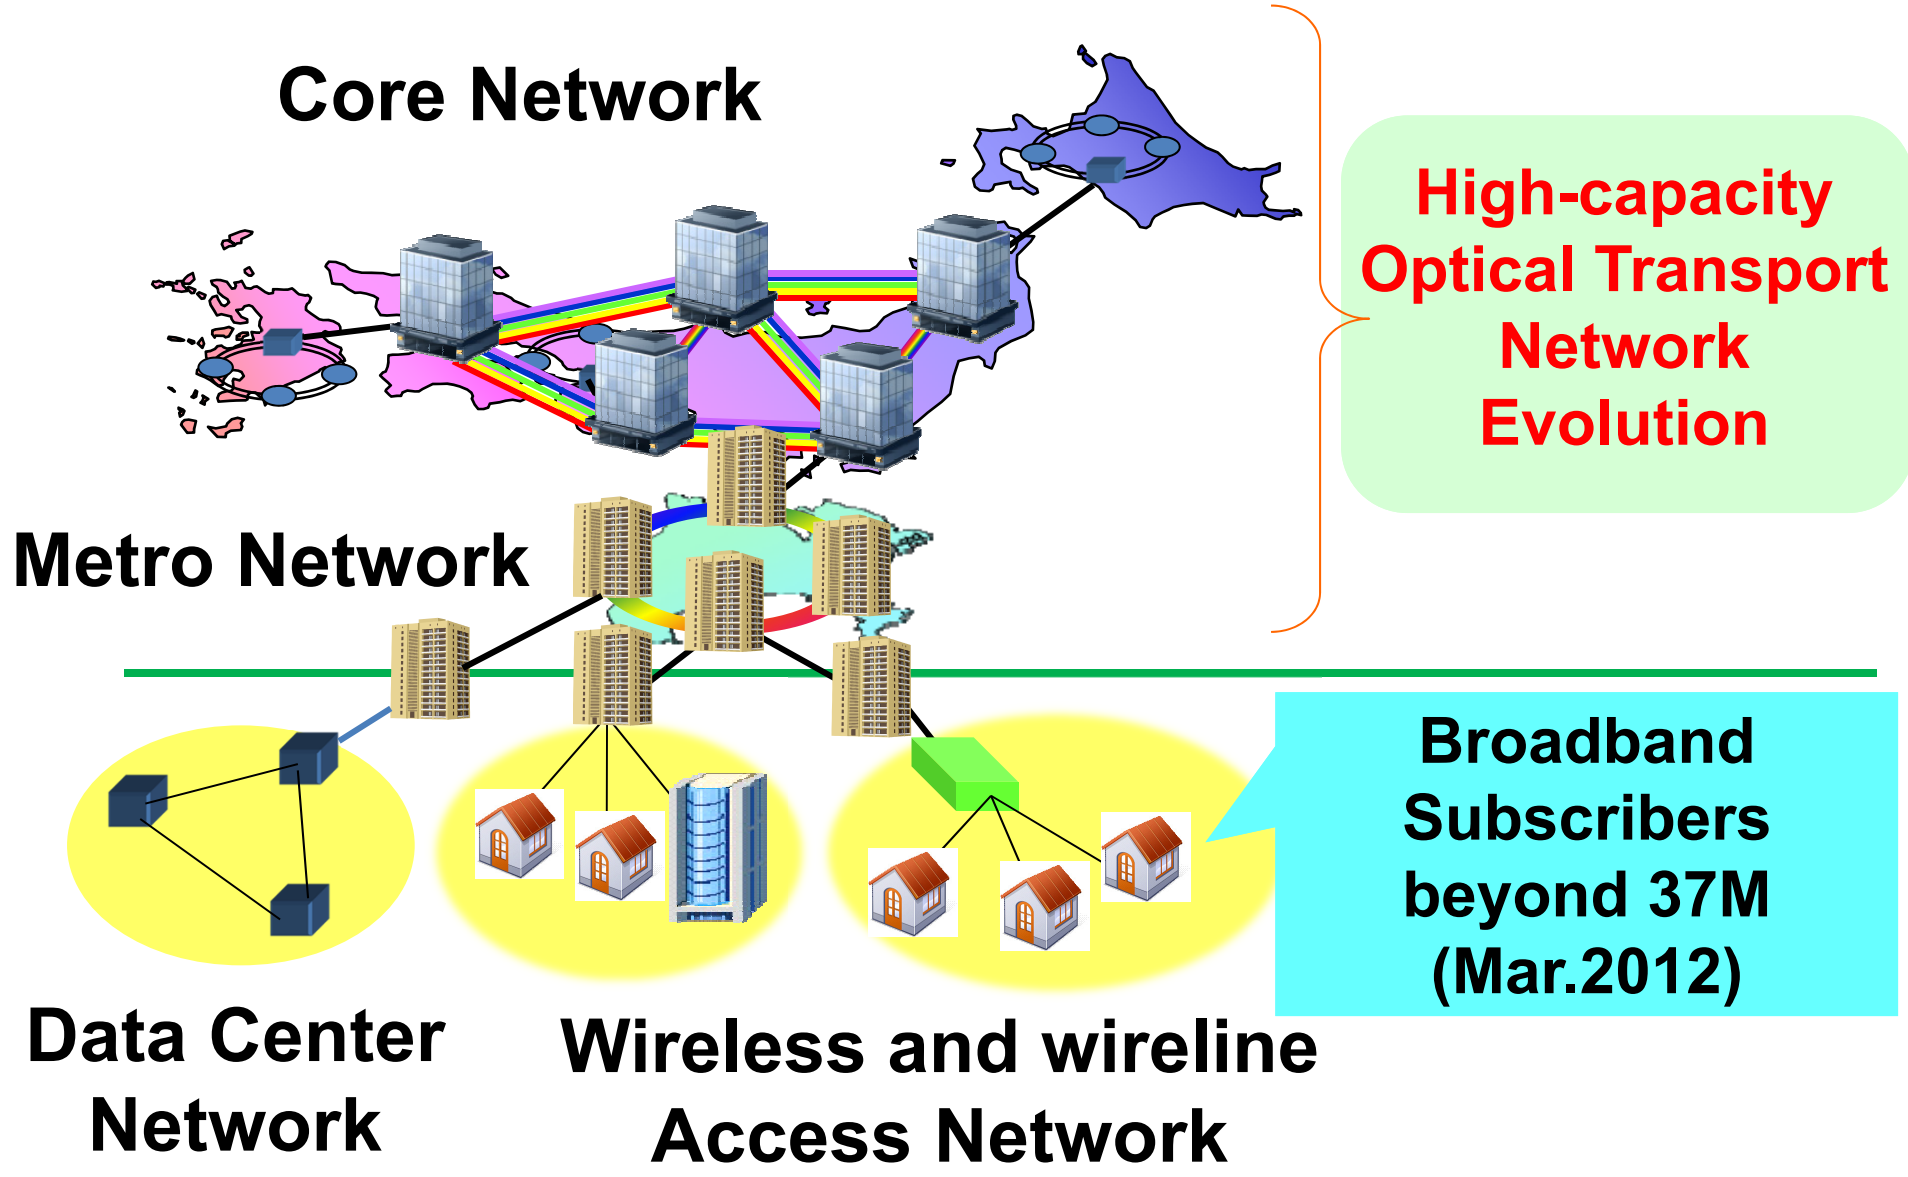

Capacity Trend in Optical Transport Systems

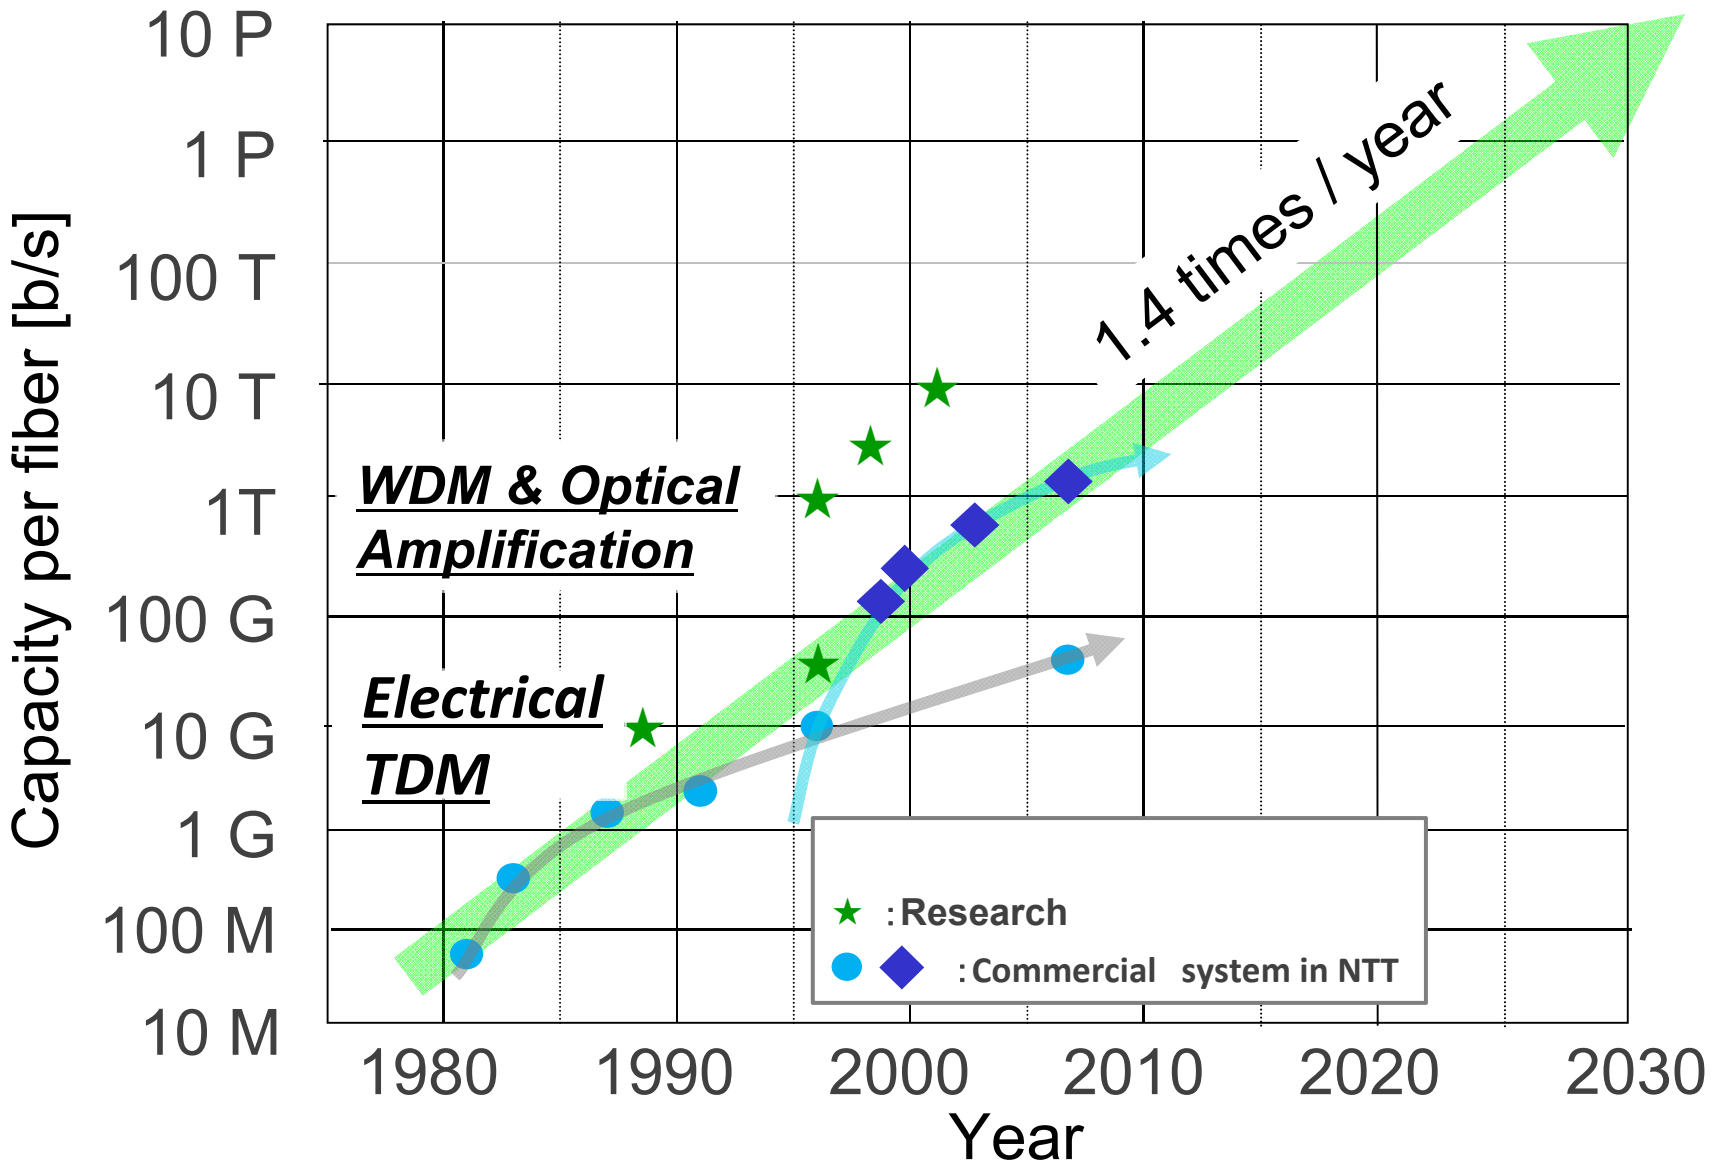

Today's Optical Transport Network

- Total Capacity = 1 Tbit/s = 100 channels x 10 Gbit/s
  (Spectral Efficiency (SE) = 0.4 bit/s/Hz)
- Modulation: 10G On-Off Keying / 40 Gbit/s DQPSK
- Typical WDM spacing: 50 / 100 GHz
  over 4 THz bandwidth at 1.5μm wavelength

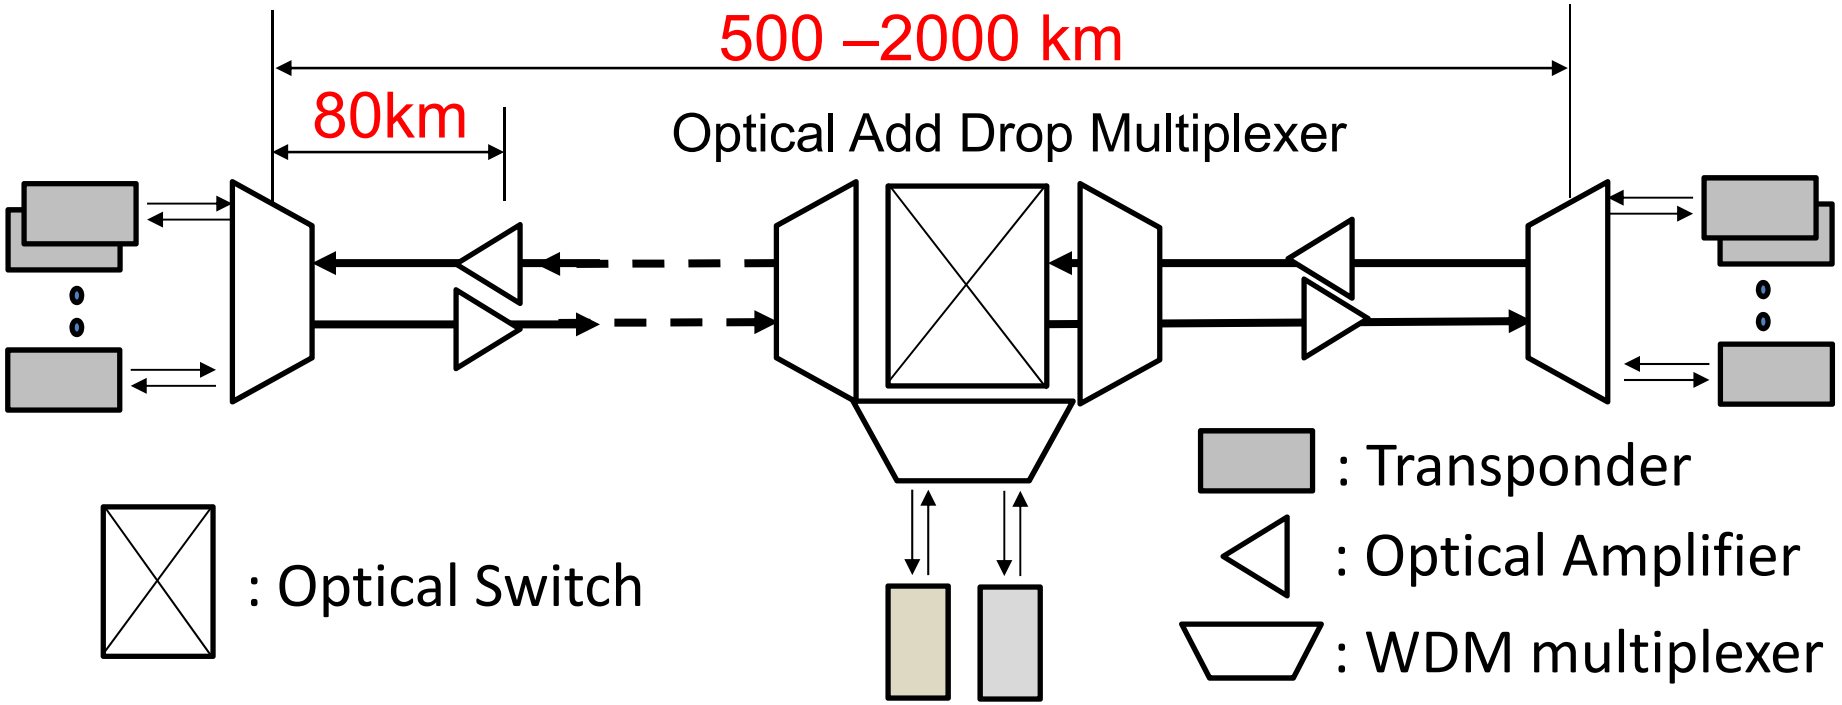

Today's Optical Transport Network

Keeping backward compatibility (repeater span, regenerative repeater span etc.) is important for flexible NW provisioning and capacity scaling

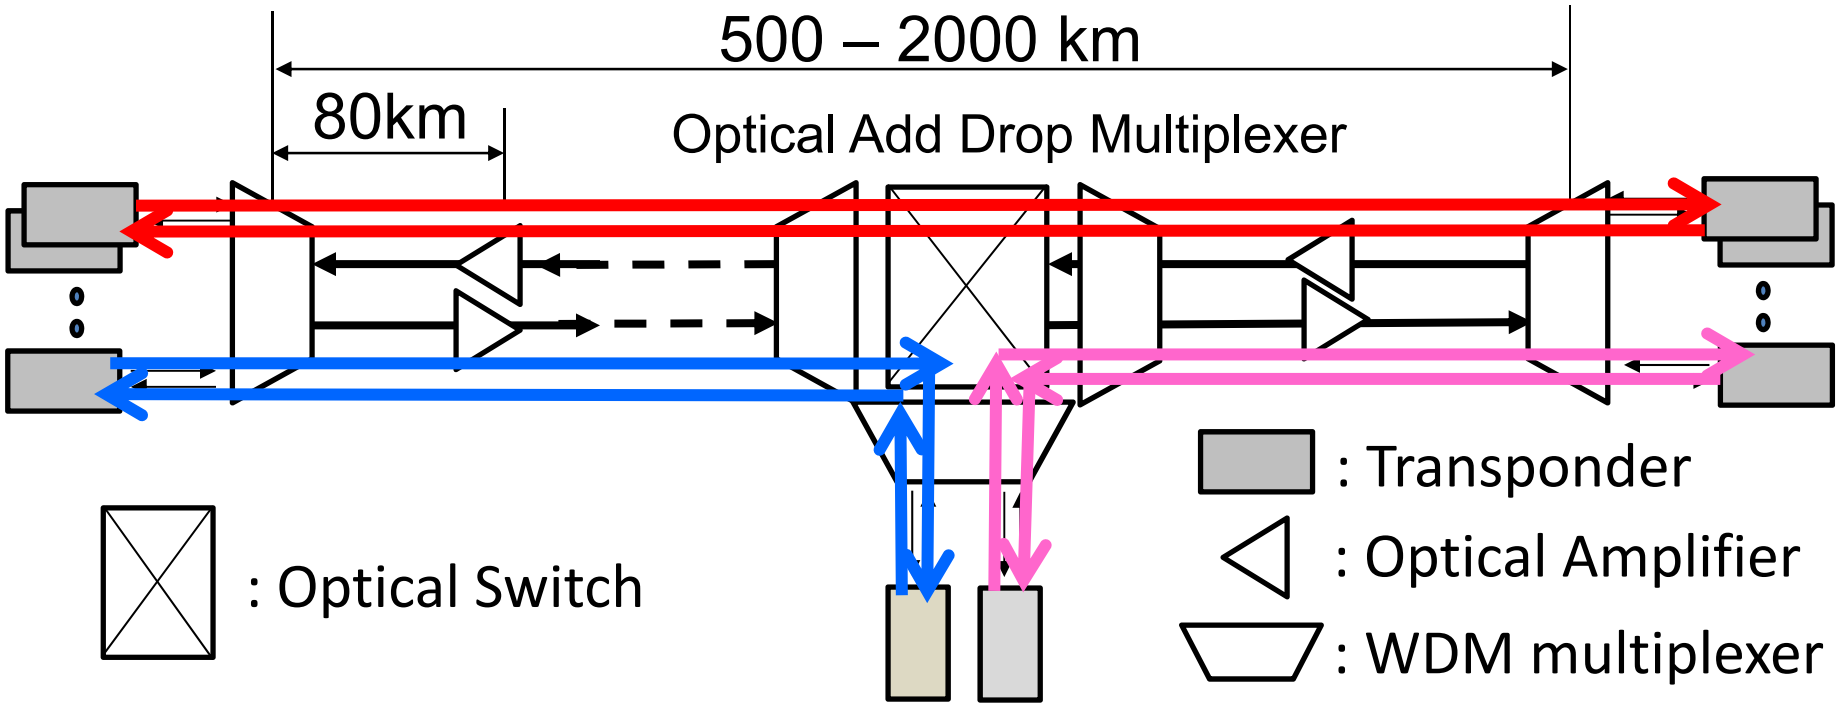

Limiting Factors of Transmission Distance

- OSNR: Optical Signal to Noise Ratio

- CD: Chromatic Dispersion
  40,000 ps/nm for 2,000km fiber
- PMD: Polarization Mode Dispersion
  100 ps; Differential Group Delay

- Fiber Nonlinearity due to Kerr effect
- Passband Narrowing by
  Optical Add-Drop Multiplexing Node

Polarization Mode Dispersion (PMD) Mitigation in High-Speed Transmission over 40 Gb/s

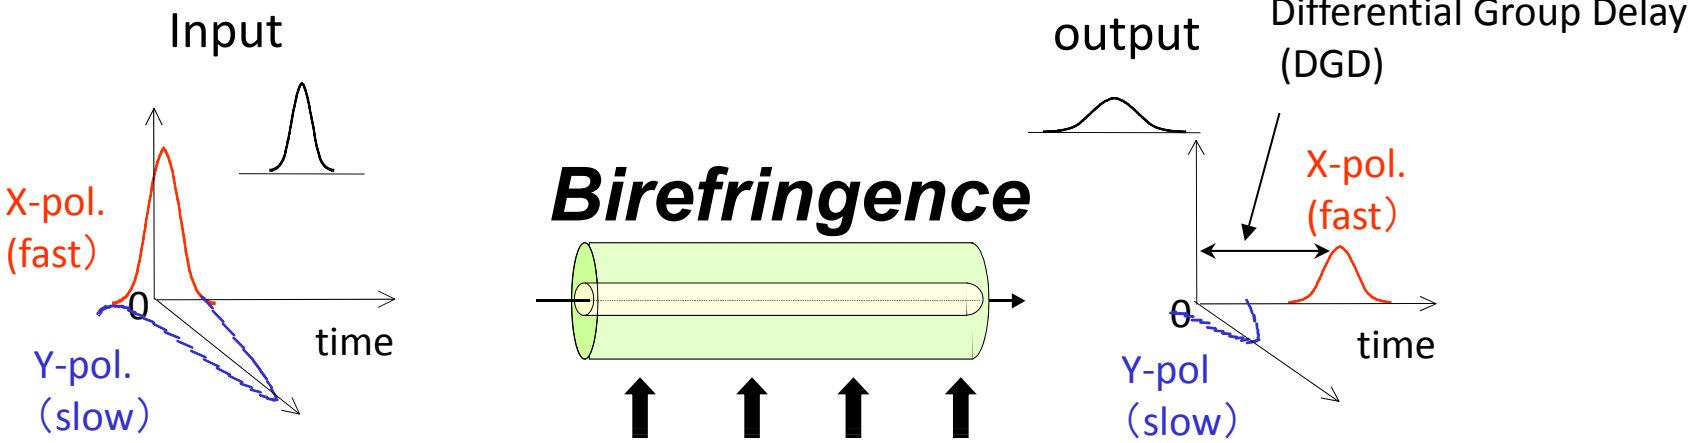

Fabrication
Installation
NW Operation

Dynamic Change
(kHz – MHz)

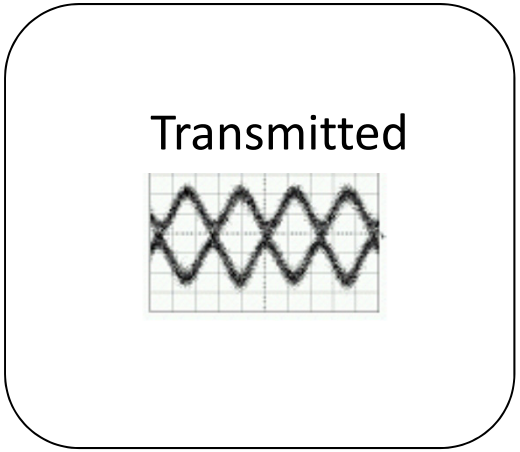

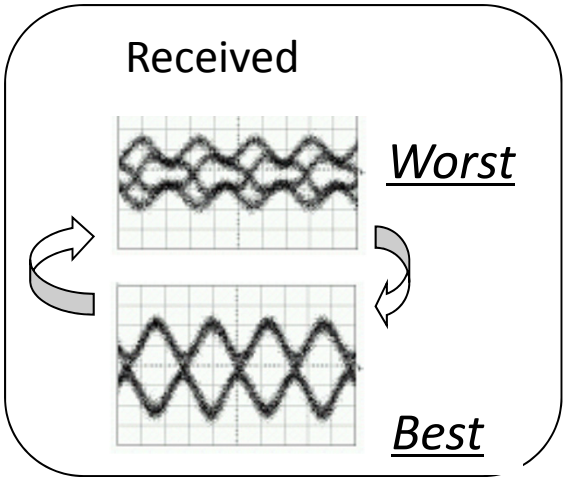

- **Reduction of symbol rate by Multi-level coding**
- **Very high-speed PMD compensation**

CD/PMD mitigation by Digital Coherent System

Digital coherent technologies are indispensable for 100G Transport.

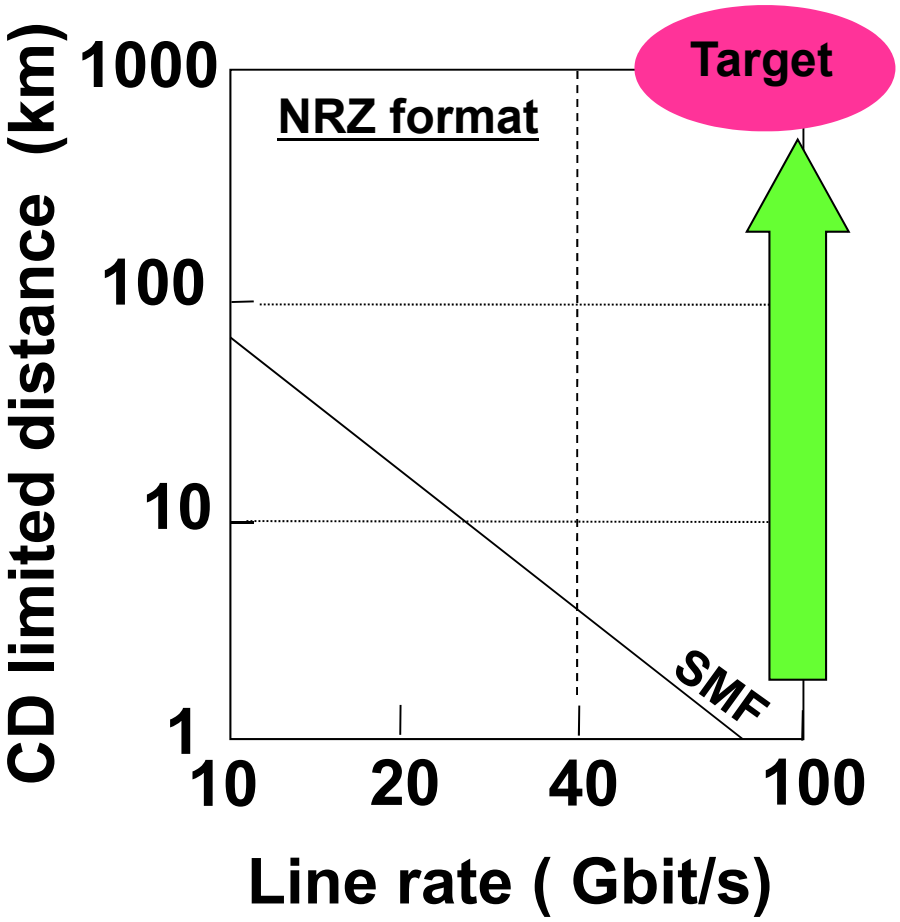

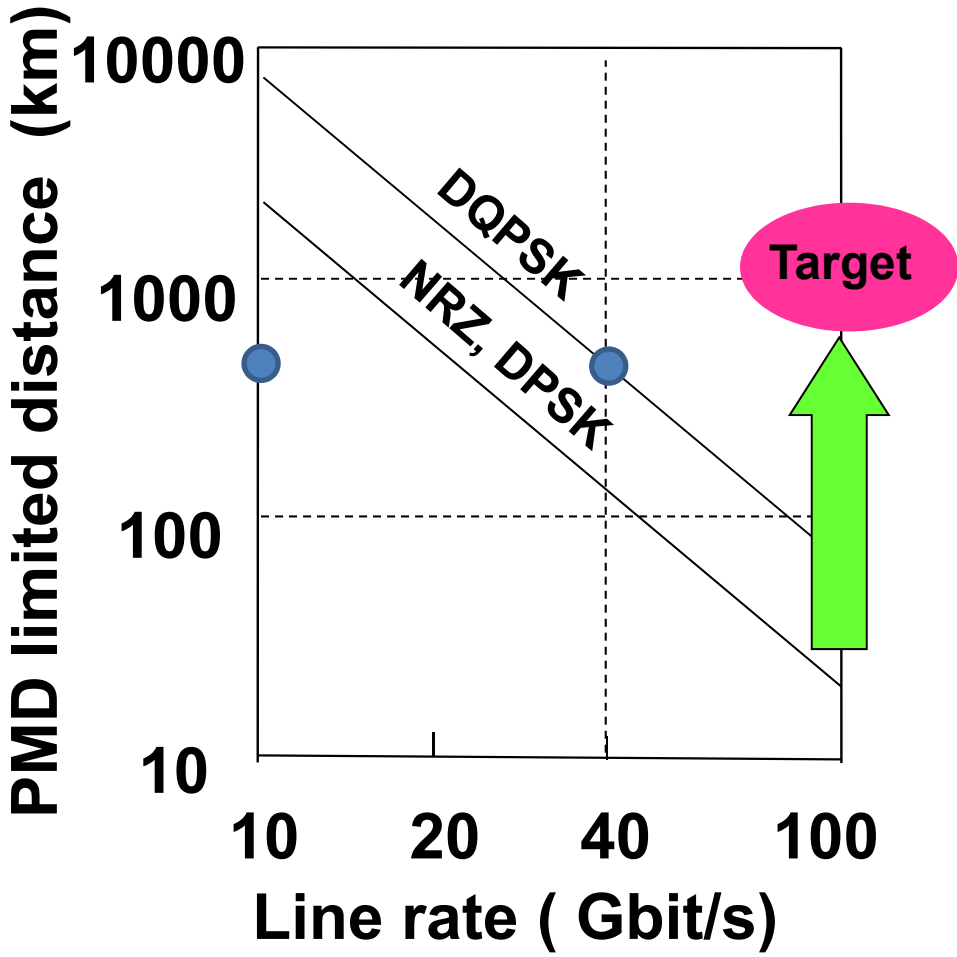

CD: Chromatic Dispersion, SMF: Single Mode Fiber
PMD: Polarization Mode Dispersion, NRZ: Non return to zero,
D(Q)PSK: Differential (Quadrature) Phase Shift Keying

100Gbit/s Digital Coherent Transport System

- **High capacity: 8-10 Tbps per fiber by enhancement of spectral efficiency**
- **Long-haul: >1000km transmission by high-sensitivity coherent detection**

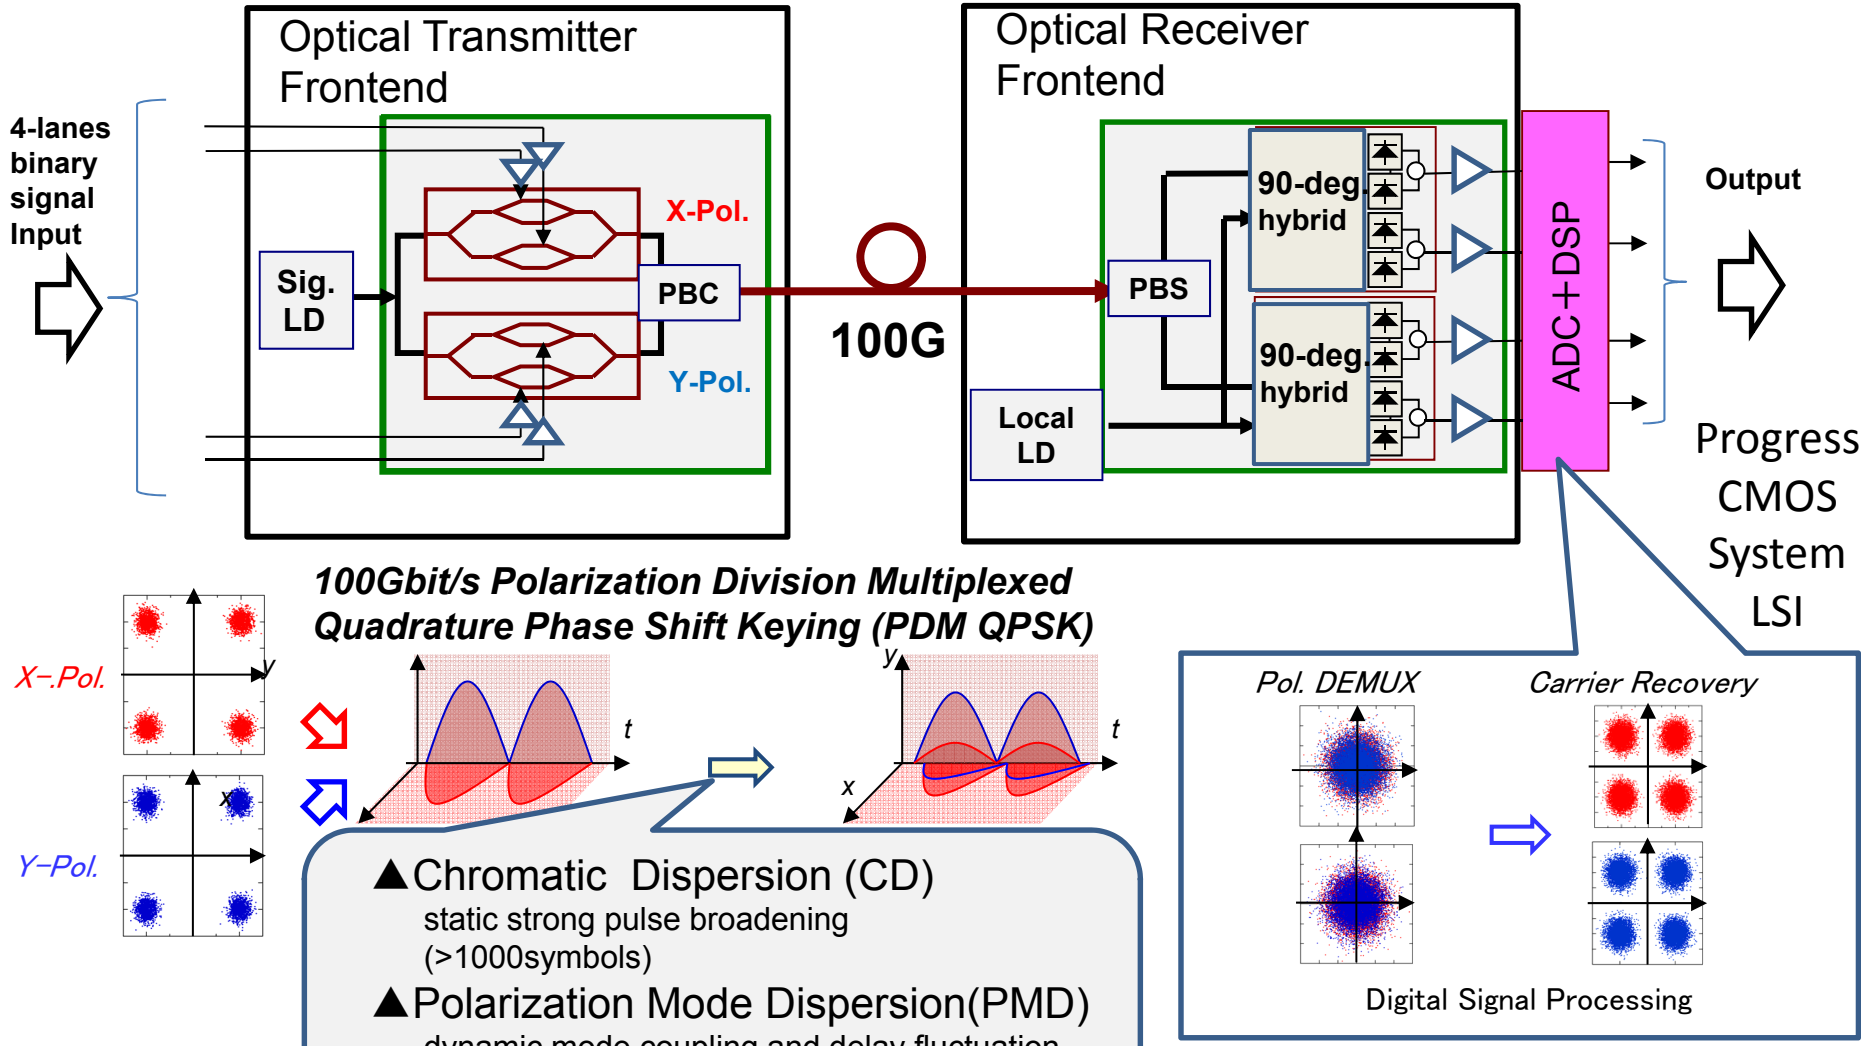

100G Digital Coherent Transport System

- OSNR: Optical Signal to Noise Ratio

  High-Net-Coding-Gain Forward Error Correction
  (Soft-Decision FEC with 20 % overhead)

- CD: Chromatic Dispersion

  Digital Fixed Equalizer by Frequency Domain Equalization

- PMD: Polarization Mode Dispersion

  Digital Adaptive Equalizer by Time Domain Equalization

- Fiber Nonlinearity

- Passband Narrowing

DSP ASIC for 100Gbit/s Digital Coherent Transport

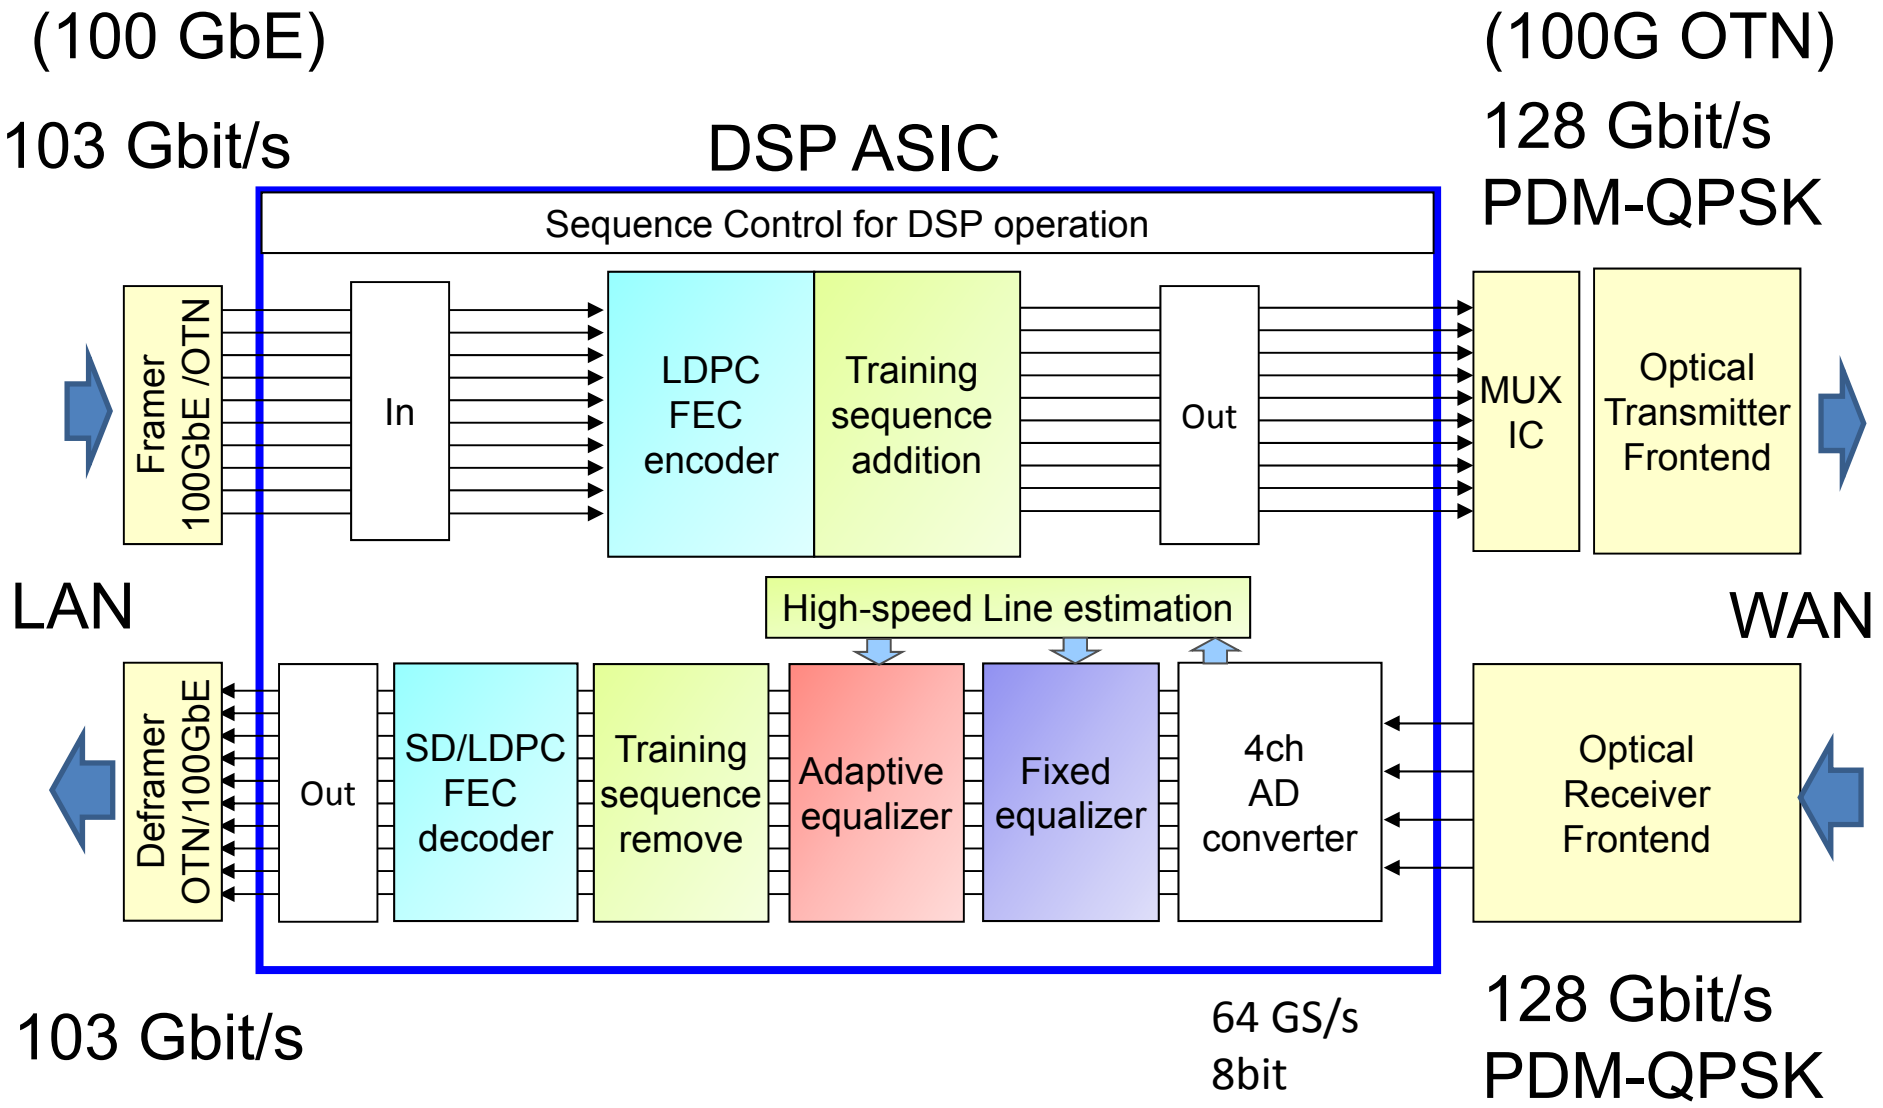

Digital Equalization in 100G Digital Coherent Transport

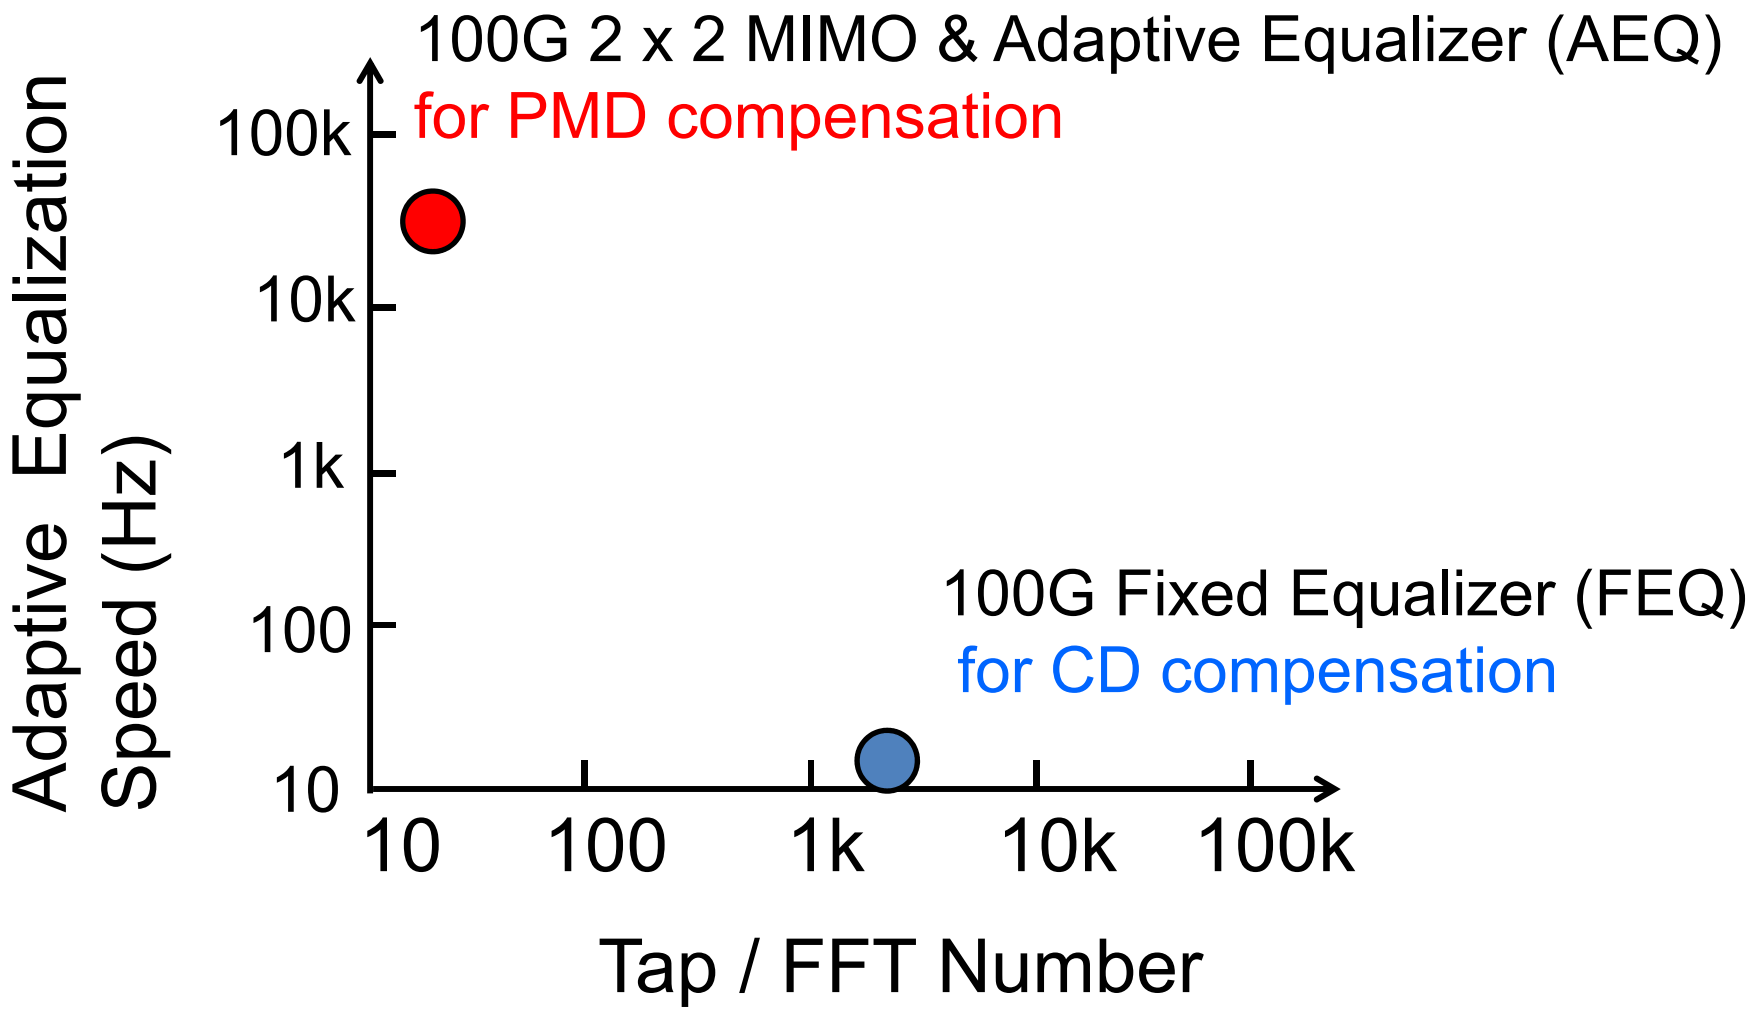

Specification and Performance of 100G DSP ASIC

Item	Specifications
Modulation format	PDM-QPSK
Line side signal bit rate	127.156 Gbps
Oversampling ratio	2.0
CD compensation	±40,000 ps/nm
Differential group delay (Polarization mode dispersion)	100 ps
Polarization tracking speed	50 kHz
Frequency offset compensation	± 5 GHz
FEC-NCG	10.8 dB @10^{-15} SD-LDPC + Enhanced-FEC
Recovery time	< 50 ms
Process	40nm CMOS

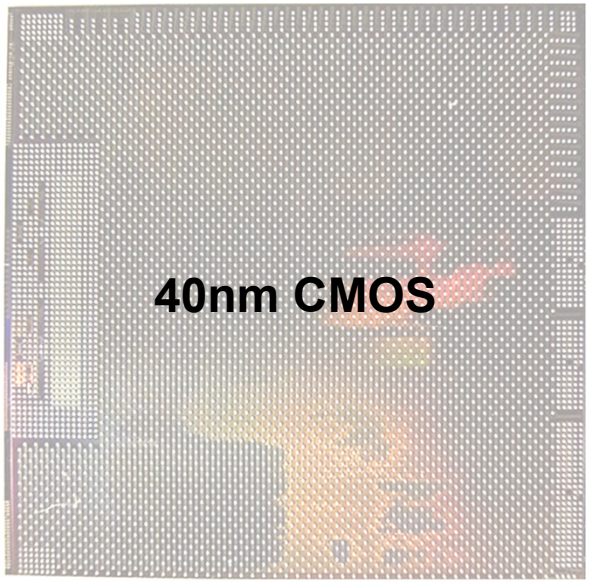

E. Yamazaki et al., Optics Express, Vol. 19, Issue 14, pp. 13179-13184 (2011)

Proof of Concept: 100Gbps Realtime DSP in 8 Tbit/s DWDM Field Experiment

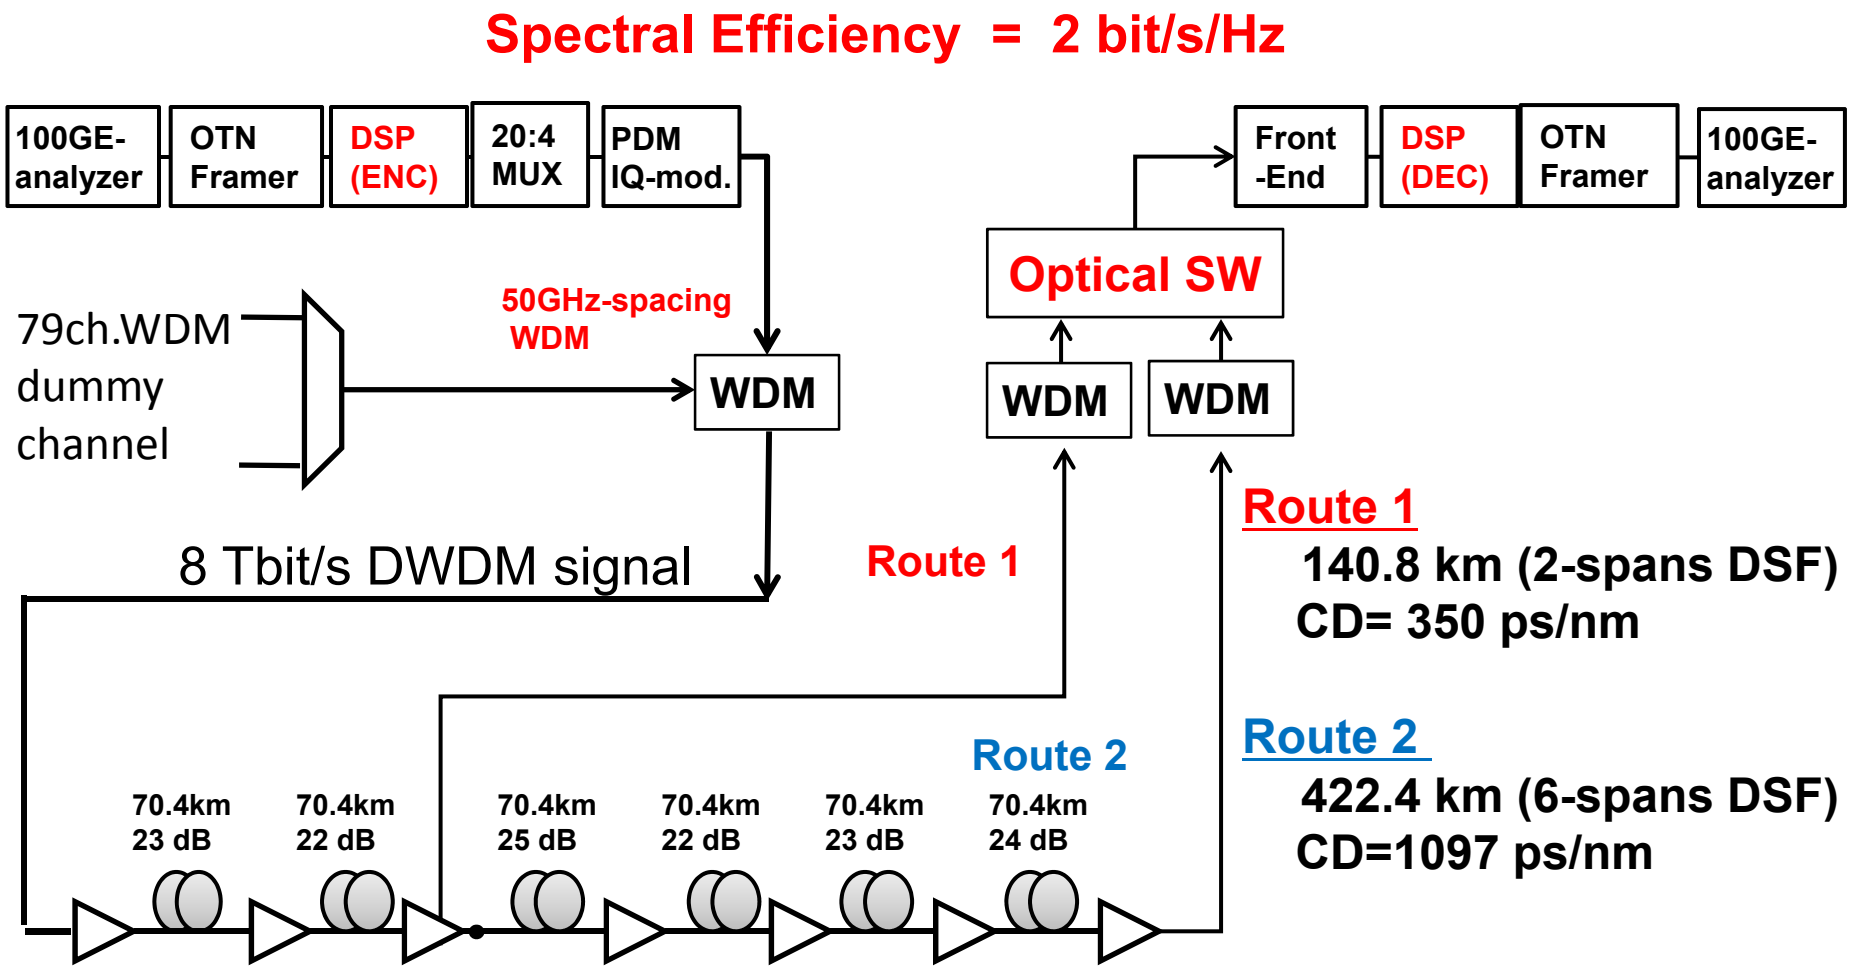

E. Yamazaki et al., Optics Express, Vol. 19, Issue 14, pp. 13179-13184 (2011)

Fast Optical Channel Recovery by 100Gbps realtime DSP

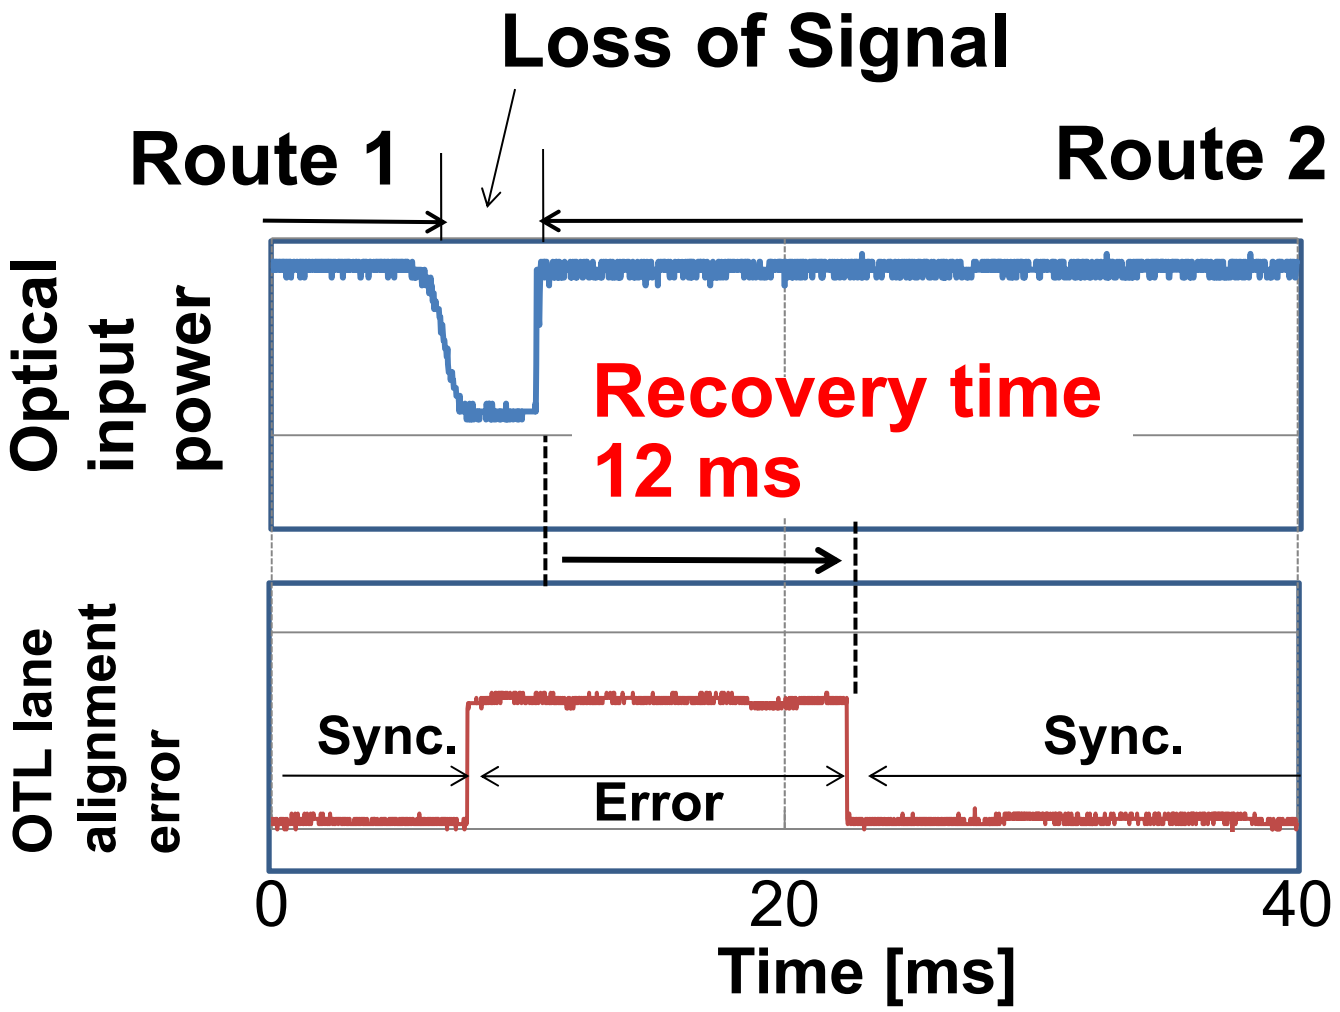

All DSP operations within 50 ms (Standard NW operation)

Next Generation Digital Coherent Transport System

Introduction of DSP-ASIC with Digital/Analog Converter into transmitter

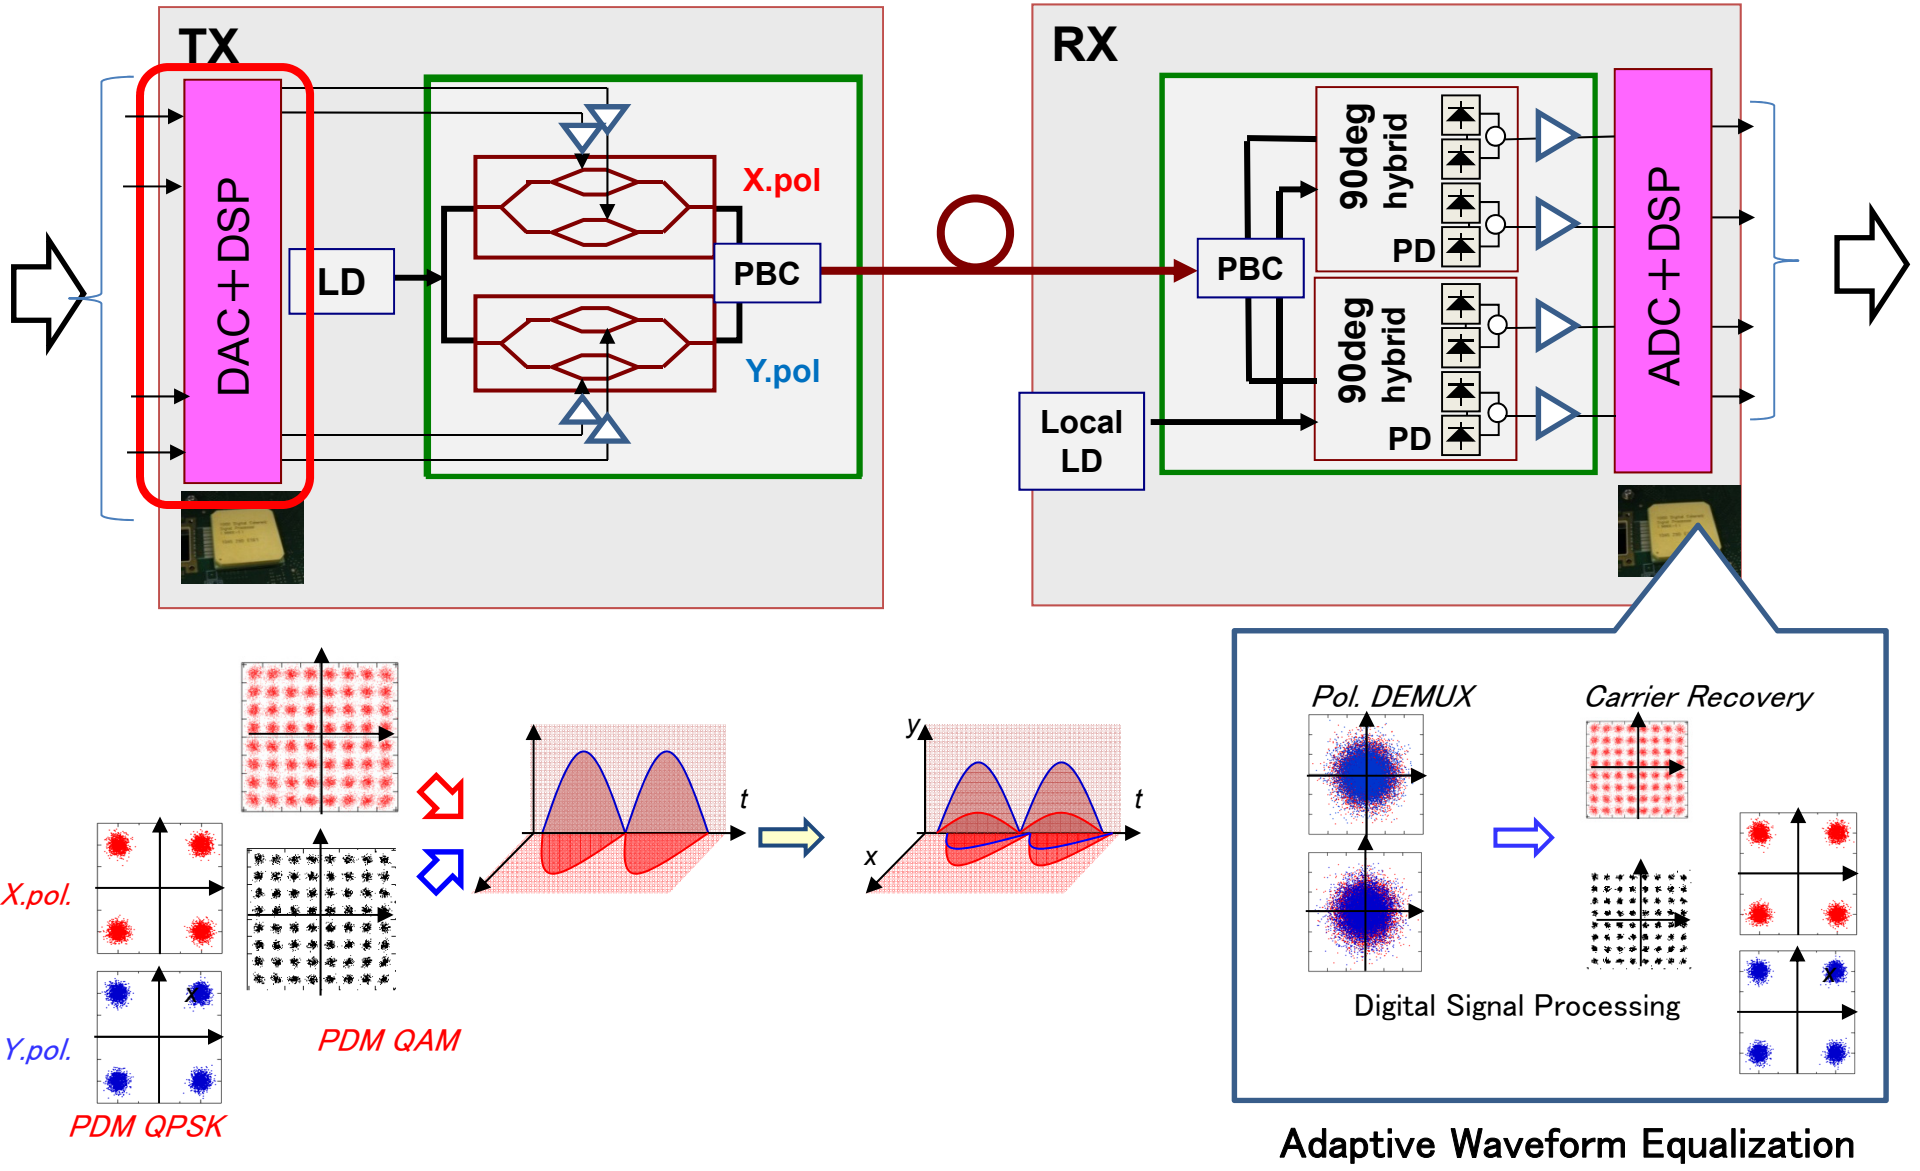

Digital Coherent Transport System Beyond 100G

- OSNR: Optical Signal to Noise Ratio
- CD: Chromatic Dispersion
- PMD: Polarization Mode Dispersion
- Fiber Nonlinearity
  - Enhancement of channel power by Digital Nonlinearity Compensation
- Passband Narrowing
  - Enhancement of Spectral Efficiency by Higher-order adaptive QAM modulation
  - Enhancement of No. of OADM by Nyquist Filtering of Transmitted Spectrum

Introduction of Multi-Carrier System for Beyond-100G (400G/1T) transmission

Beyond-100G/ channel system: Multi-carrier system

400G = 4 x 100G PDM-QPSK Sub carriers ,
= 2 x 200G PDM-16QAM Sub carriers etc.

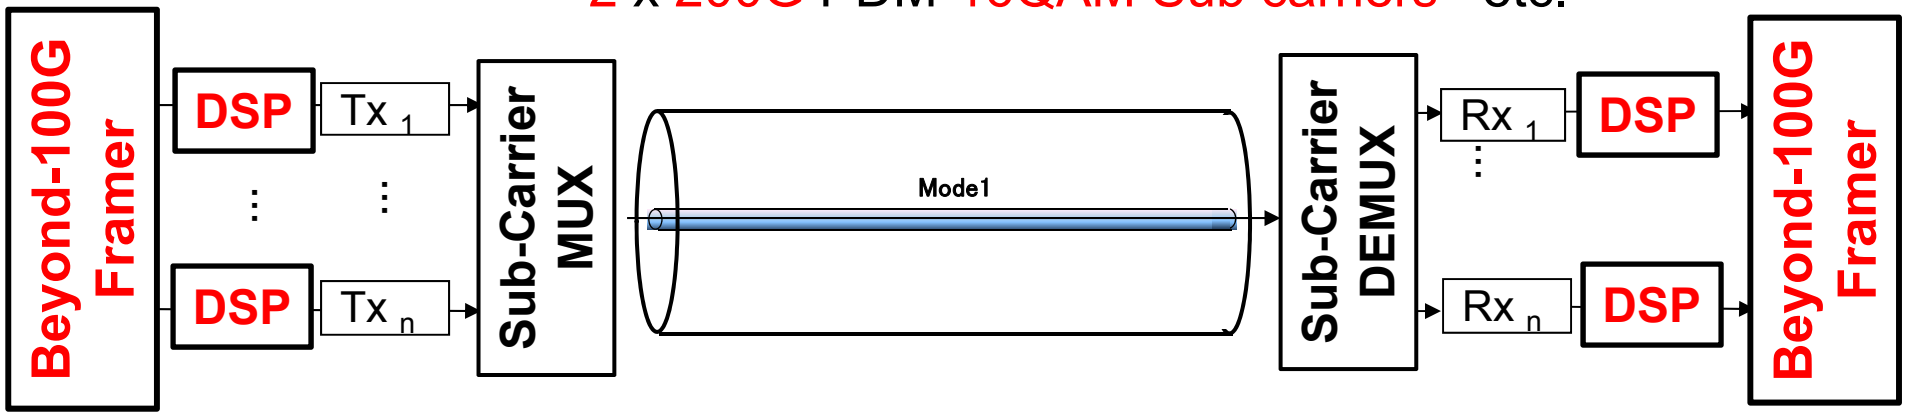

Flexible Bandwidth Demand *ADC/DAC Bandwidth Limit*

100G channel system: Single-carrier system

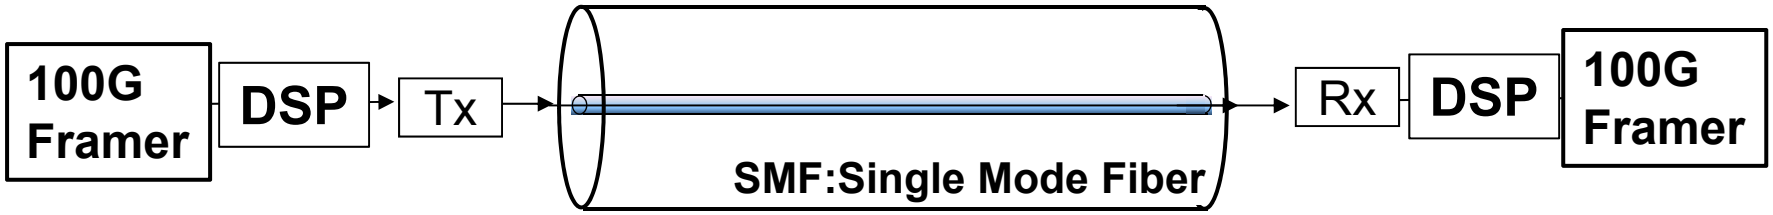

Tx: Transmitter Frontend

Rx: Receiver Frontend

Capacity Trend in Optical Transport Systems

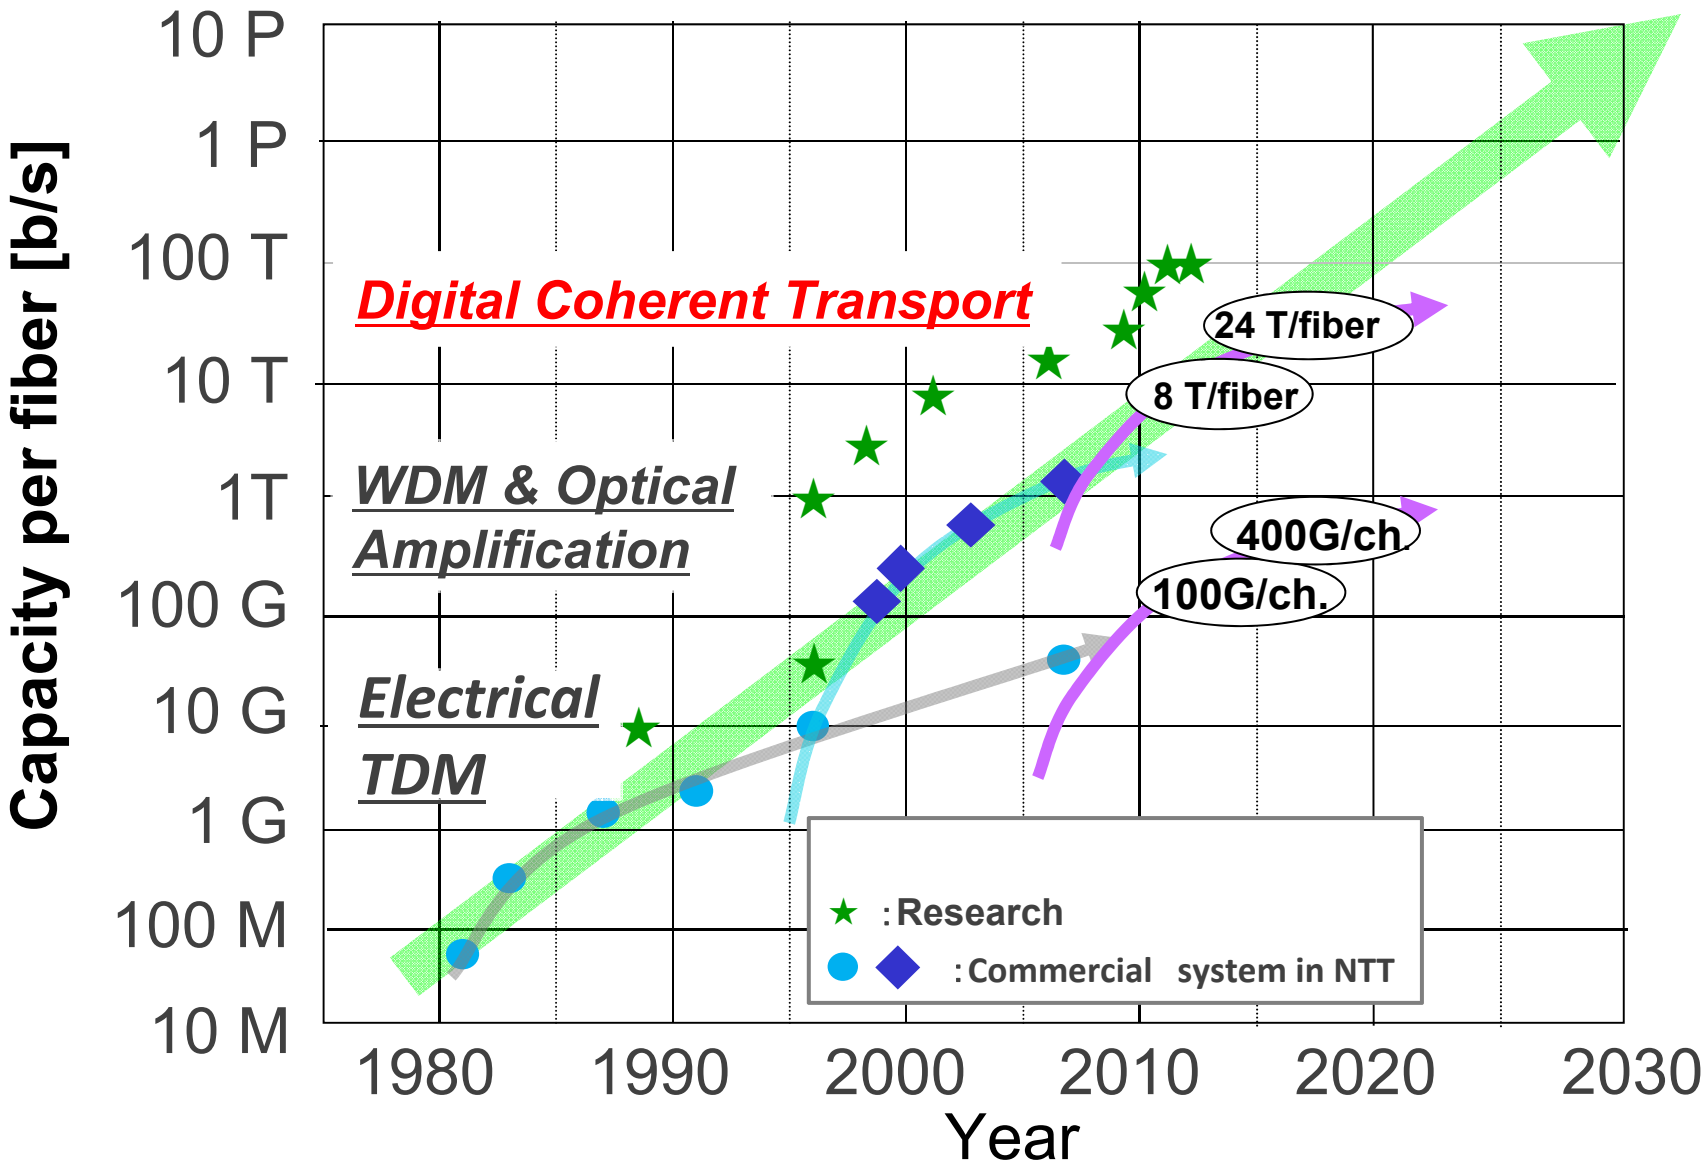

Capacity Trend in Optical Transport Systems

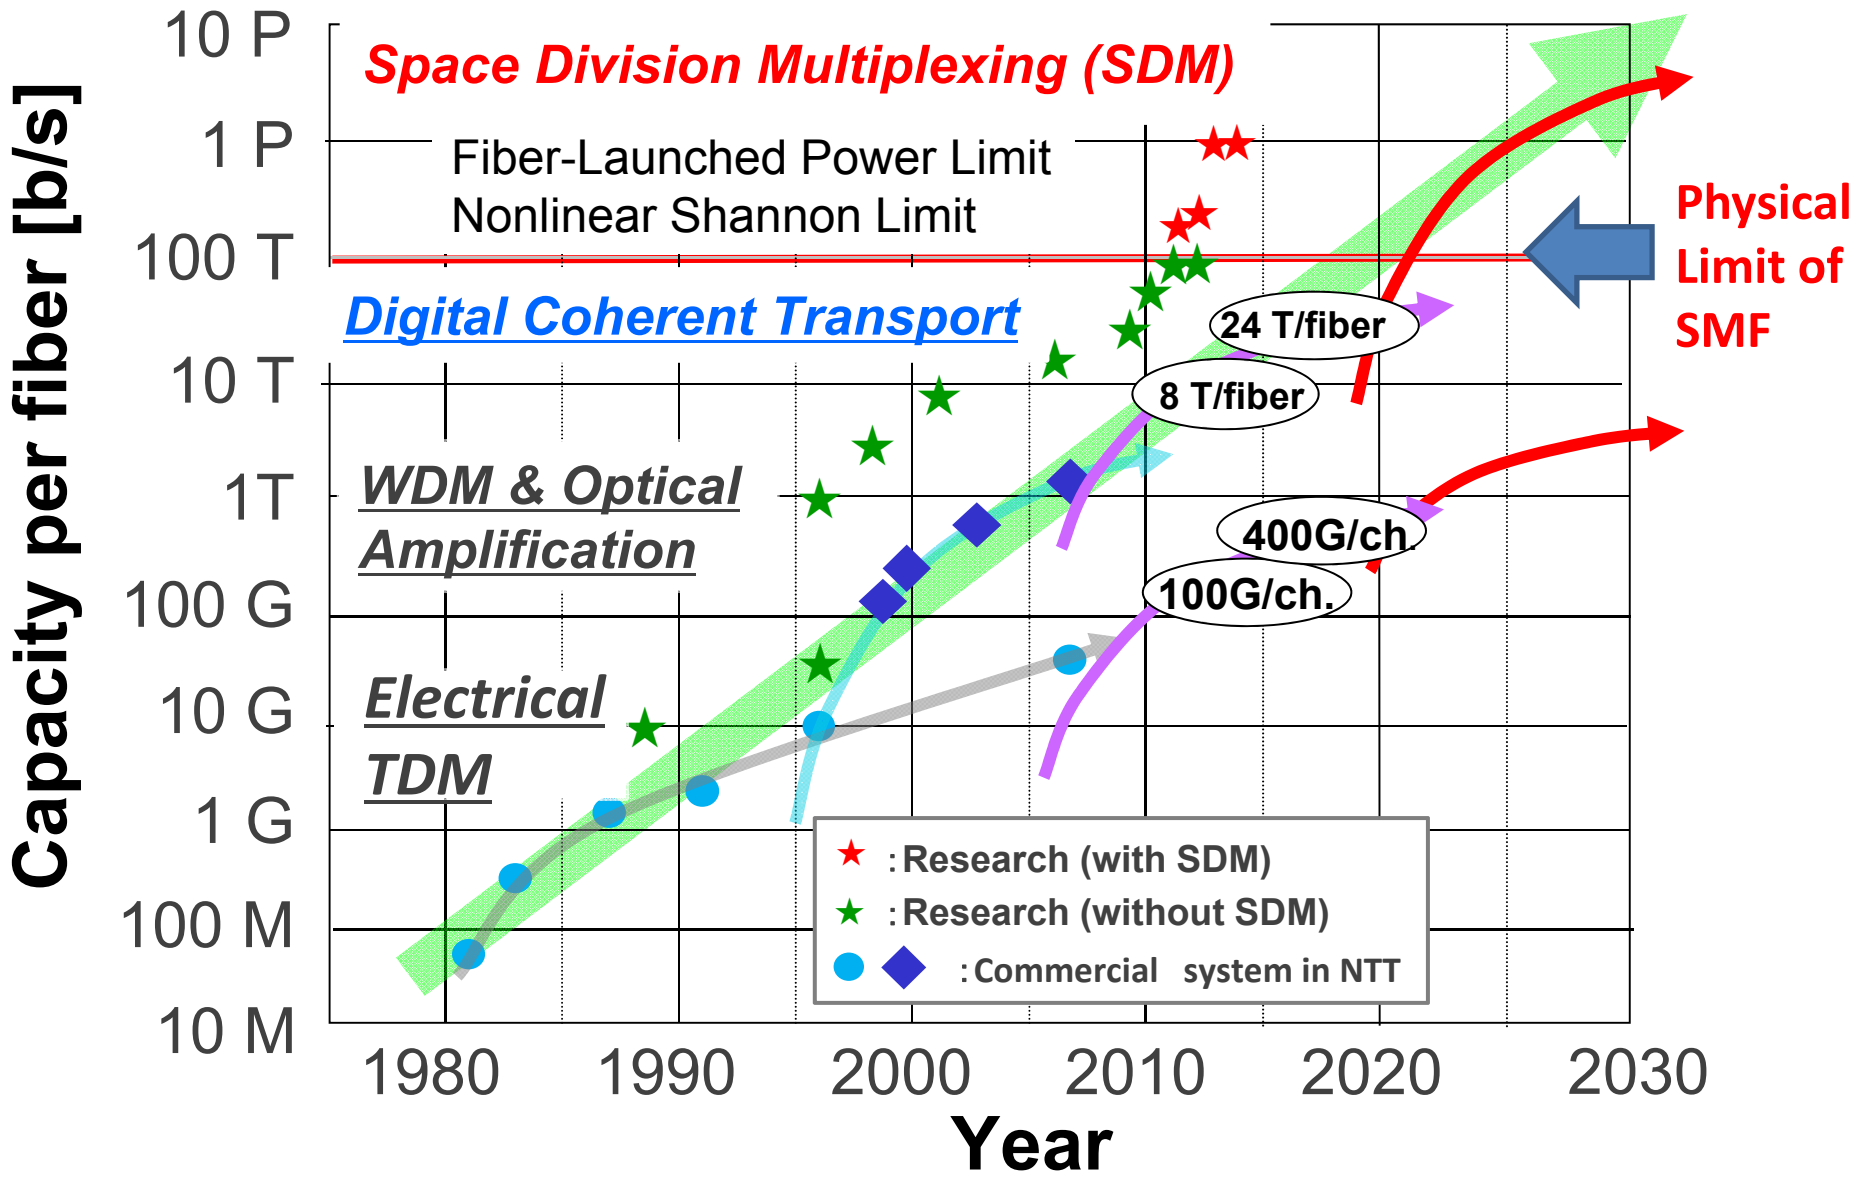

Scalable Optical Transport System based on Space Division Multiplexing (SDM)

Capacity Scaling: M x N > 100

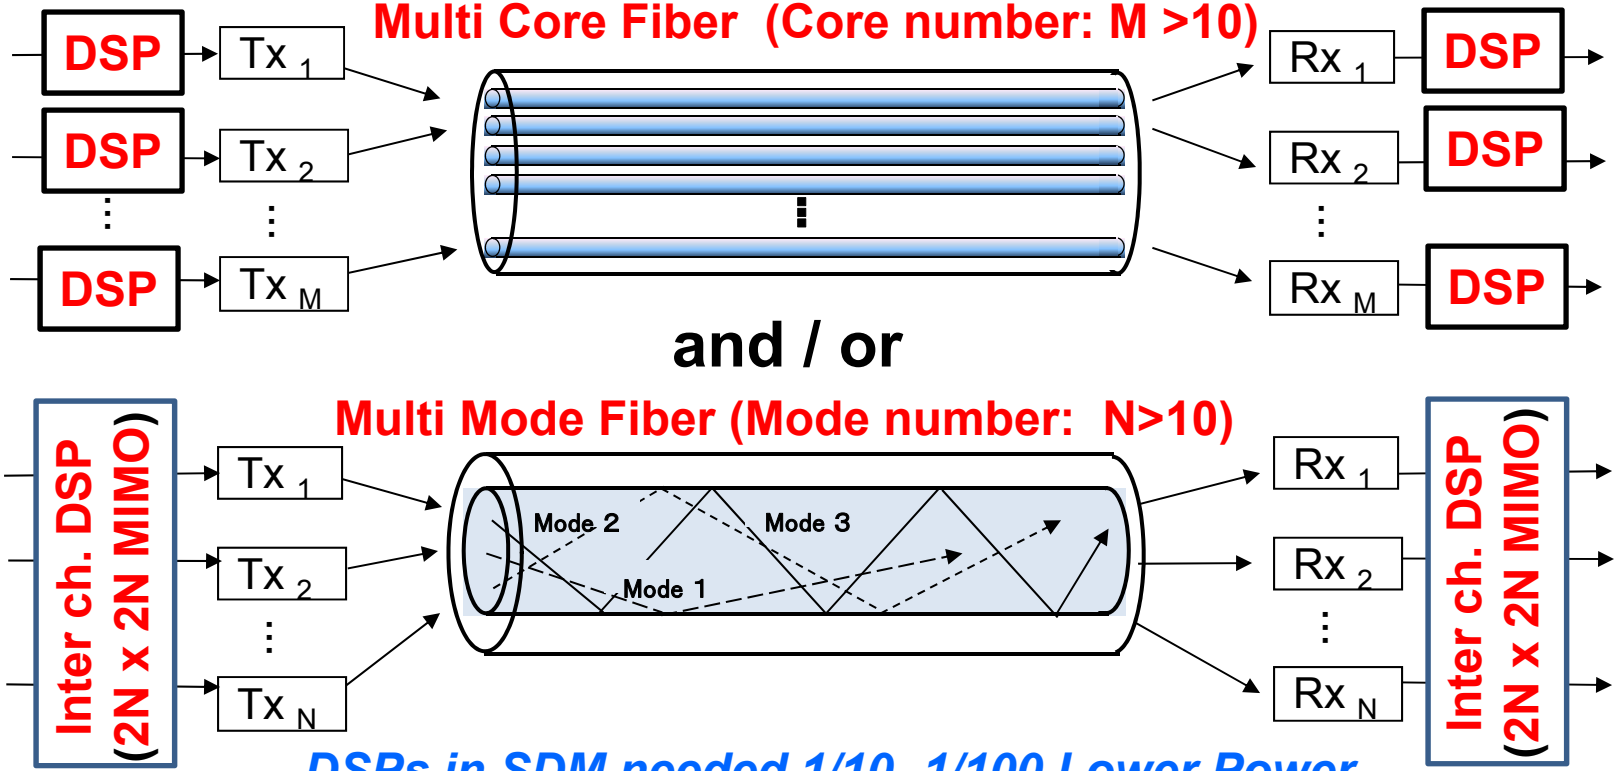

DSPs in SDM needed 1/10- 1/100 Lower Power

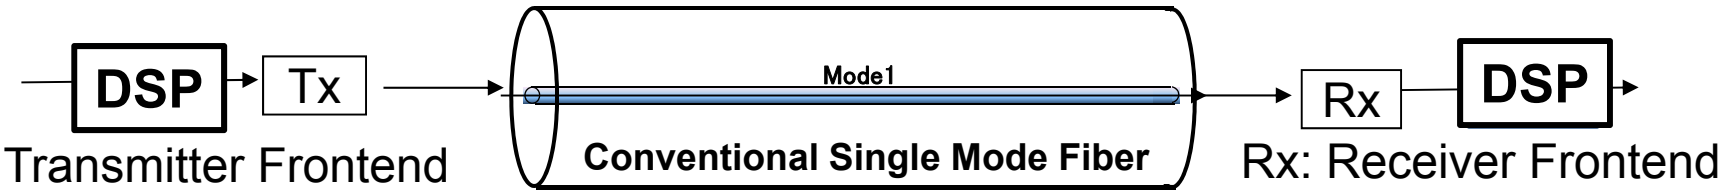

Tx: Transmitter Frontend

Rx: Receiver Frontend

Transmission Distance Limit in SDM

- OSNR: Optical Signal to Noise Ratio
- CD: Chromatic Dispersion （20 ps/nm/km)
- PMD: Polarization Mode Dispersion
- Fiber Nonlinearity due to Kerr effect
- Passband Narrowing by Optical Add-Drop Node

- Crosstalk / Modal Dispersion Compensation
  - Inter Core Crosstalk
  - Modal Dispersion; 0.5 - 2 ns/km

Classification and features of SDM system

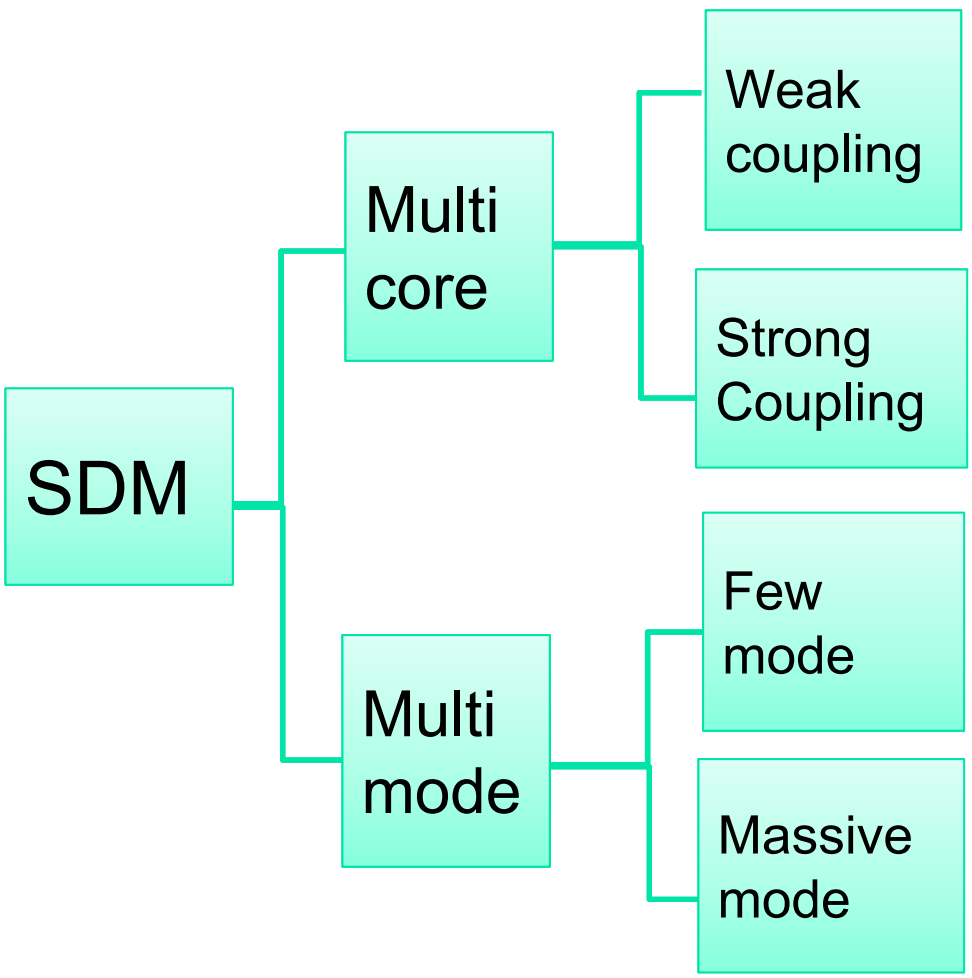

Multiplicity (Reported)	DSP complexity
19 cores	- 2 x 2 MIMO - Ch.by Ch DSP
3 cores	
6 modes (excluding pol. mode)	- 12 x 12 MIMO
N mode Potentially Large	- 2N x 2N MIMO - Inter ch. DSP

Trade off between DSP Complexity and SDM multiplicity

Challenge of Pb/s-class Capacity Scaling

DSP complexity: Key Features of 400G

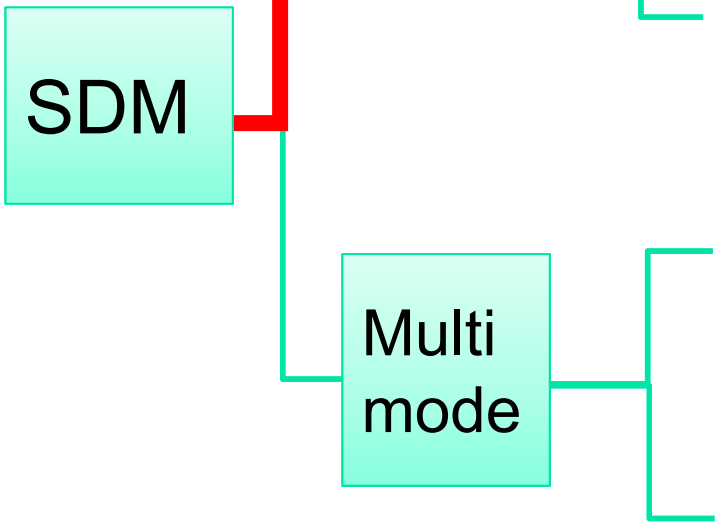

- PDM 32 QAM format
- Nyquist Filtering
- 400G with 8-subcarrier multiplexing
- 2 x 2 MIMO (for PDM)
- (20% SD-FEC)

Offline processing

First 1 Pb/s Transmission Experiment over 52km

（NTT, Fujikura, Hokkaido Univ. DTU）

1.01-Pb/s Capacity = 12 core x 84. 4 Tb/s/core

H. Takara et al, ECOC2012, Th.3.C.1. 2012

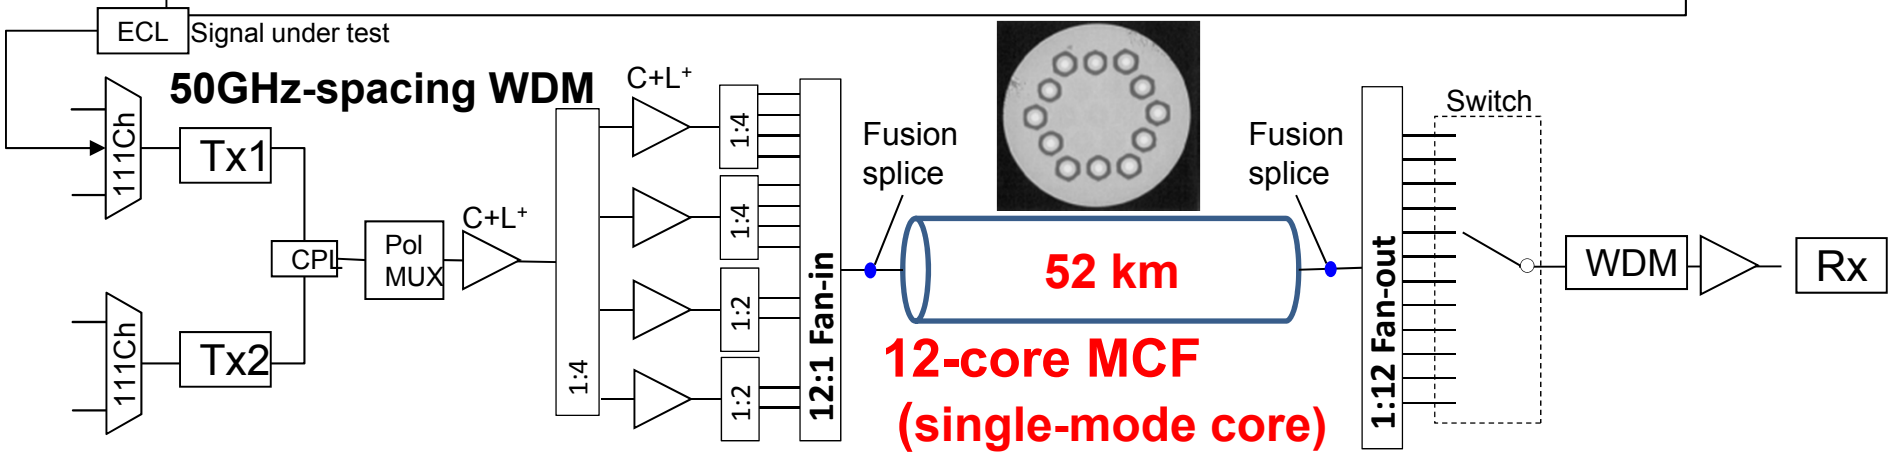

456.8Gb/s/ch. PDM 32-QAM (8-subcarrier multiplexing)

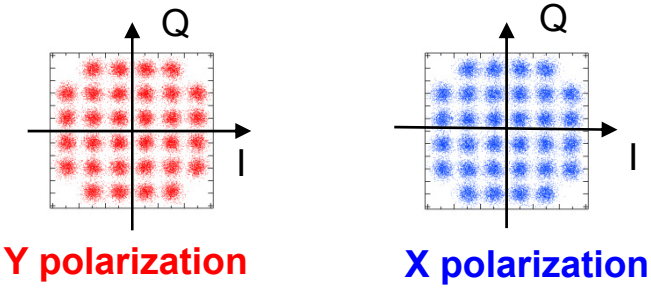

Spectral Efficiency= 7.6 b/s/Hz/fiber

84.4 Tb/s / core by 222 ch. DWDM

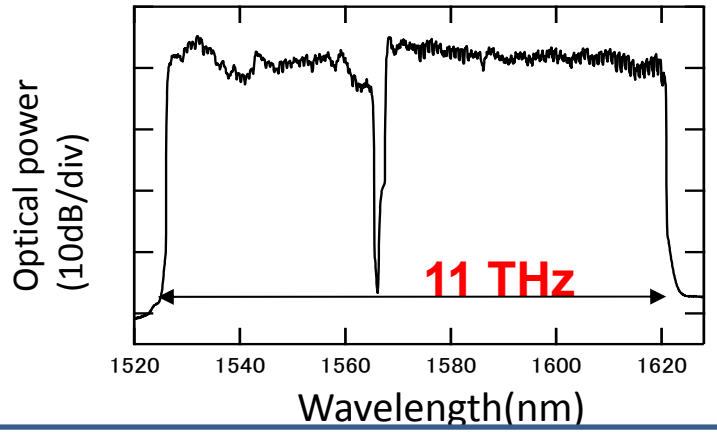

Issue for Massive MIMO DSP for Modal Dispersion Compensation in SDM; Mode-Division Multiplexing

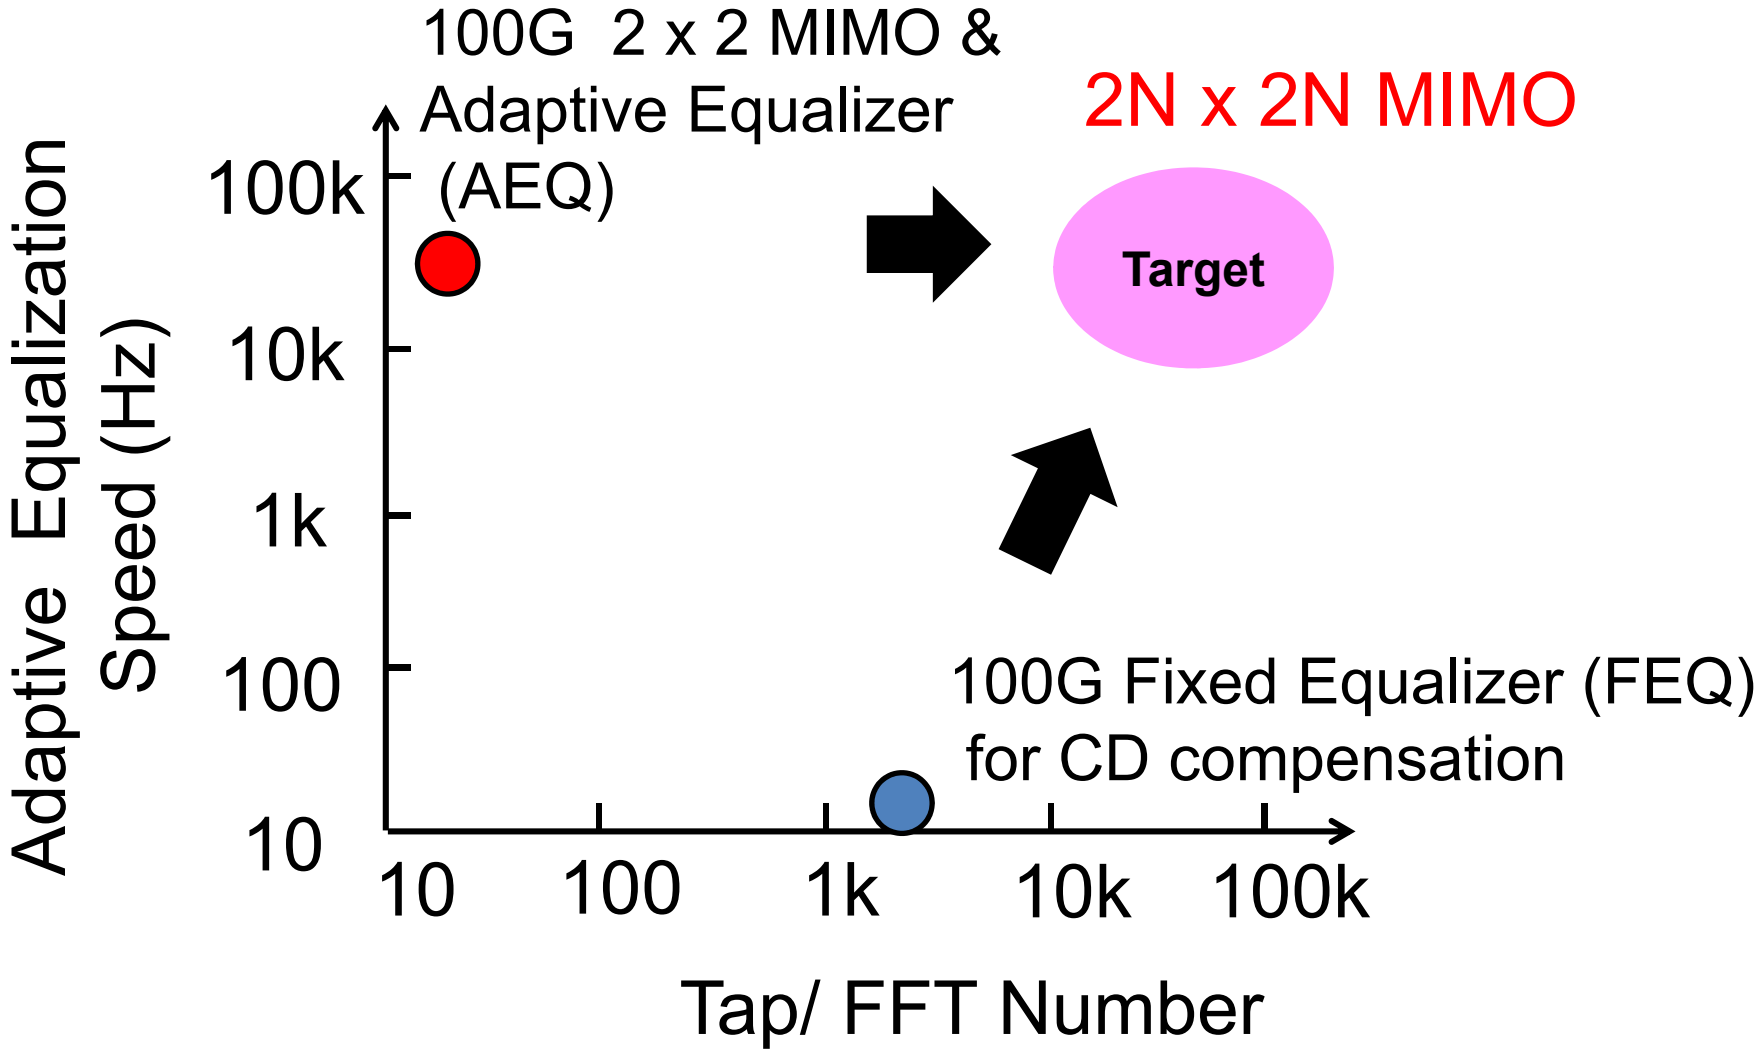

Summary

- Digital Coherent Transport System based on DSP
  - 8-Tbps capacity commercial system deployment based on 100G,
  - >20Tbps capacity transport technologies are under development beyond 100G.

- Future Pb/s-class capacity evolution based on Space Division Multiplexing (SDM)
  - Ultra-low-power DSP-ASIC architecture and design is key enabler.

Acknowledgement

- **This work is partly supported by the following R&D projects**

- MIC project “Research and Development on High Speed Optical Transport System Technologies“
- MIC project “Research and Development on Ultra-high Speed Optical Edge Node Technologies”
- MIC project “The research and development project for the ultra-high speed and green photonic networks”
- NICT project “Research and Development of Optical Transmission Technologies for transparent metro access network (Lambda-reach)”
- NICT project “Innovative Optical Communication Infrastructure“
- NICT project “Innovative Optical Fiber Technologies“

- **The authors thanks NEC Corporation, Fujitsu Limited, and Mitsubishi Electric Corporation.**

XILINX

ALL PROGRAMMABLE™

A Heterogeneous 3D-IC Consisting of Two 28nm FPGA Die and 32 Reconfigurable High-Performance Data Converters

Christophe Erdmann[1], Donnacha Lowney[1], Adrian Lynam[1], Aidan Keady[1], John McGrath[1], Edward Cullen[1], Daire Breathnach[1], Denis Keane[1], Patrick Lynch[1], Marites De La Torre[1], Ronnie De La Torre[1], Peng Lim[1], Anthony Collins[1], Brendan Farley[1], Liam Madden[2]

[1]Xilinx, Dublin, Ireland

[2]Xilinx, San Jose, CA

Outline

- **Motivation** – Integration partitioning
- **3D-IC** – Enabling technology
- **Architecture** – Block diagram, isolation strategy and analog reconfigurability
- **Measurements** – Analog performance and crosstalk
- **Summary**
- **Conclusion**

Motivation

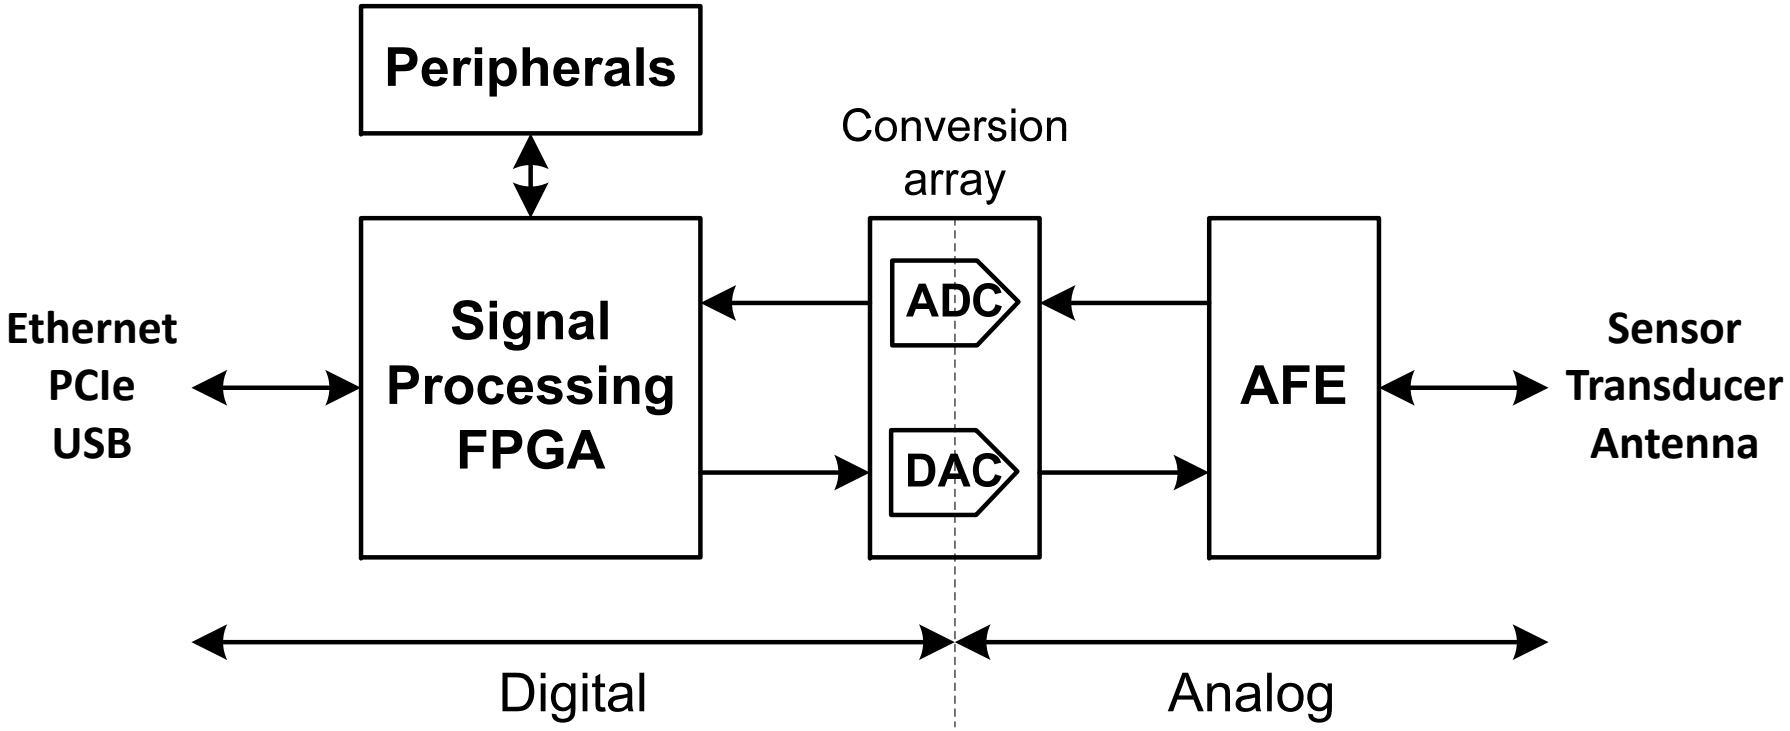

- **Typical signal processing chain includes Analog Front End, ADC/DAC array and FPGA**

Motivation

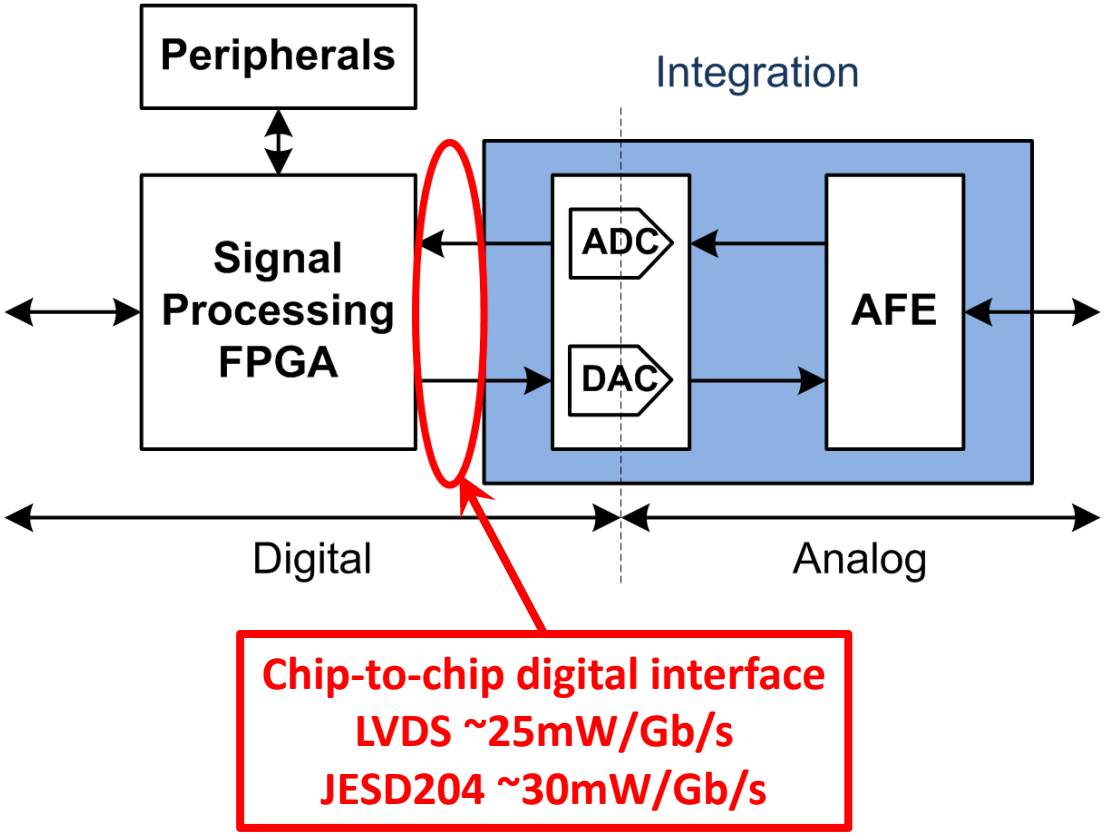

- **Integration of analog with data converters saves board space BUT does NOT reduce system power**

Motivation

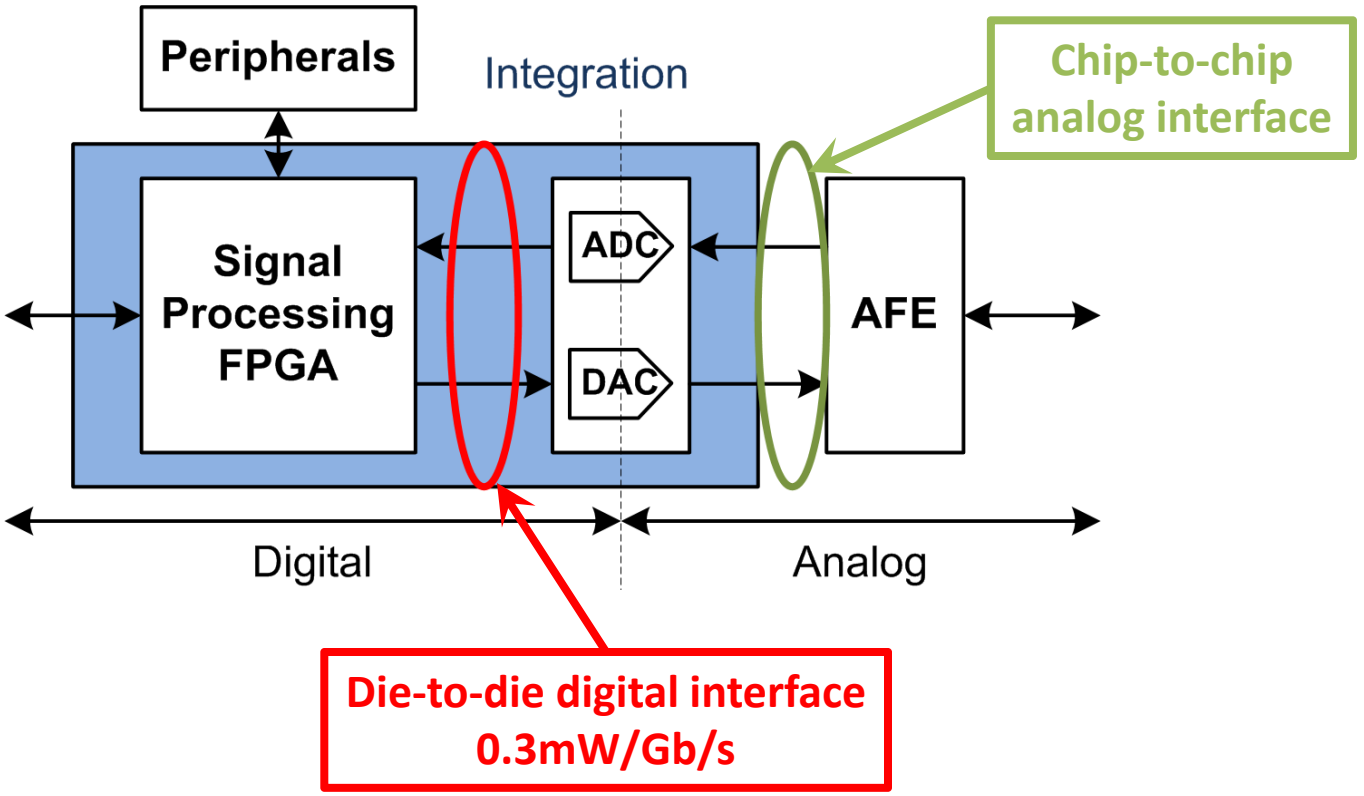

- **Integration of FPGA with data converters saves board space AND reduces system power (interface power reduced by ~100x)**

Silicon Stack Interposer technology

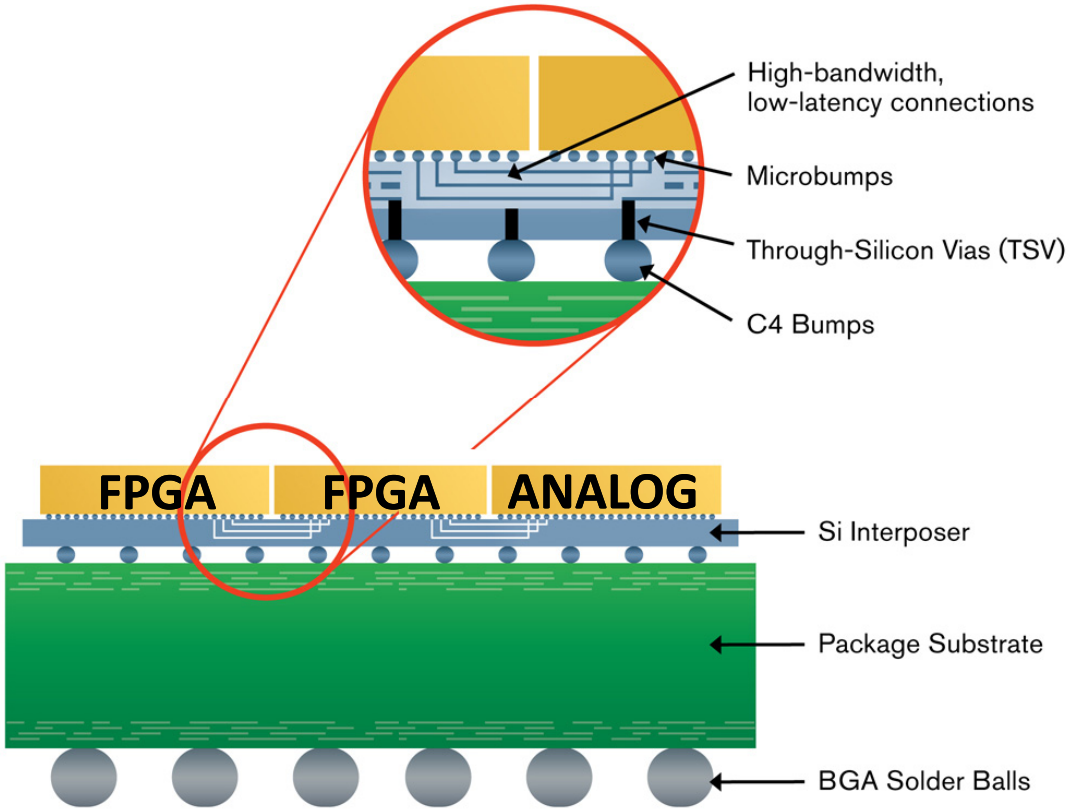

- Heterogeneous process can coexist within the same package
- No thermal or stress effect on analog matching
- Die-to-die interconnect is assured by the passive interposer

SSI cross section

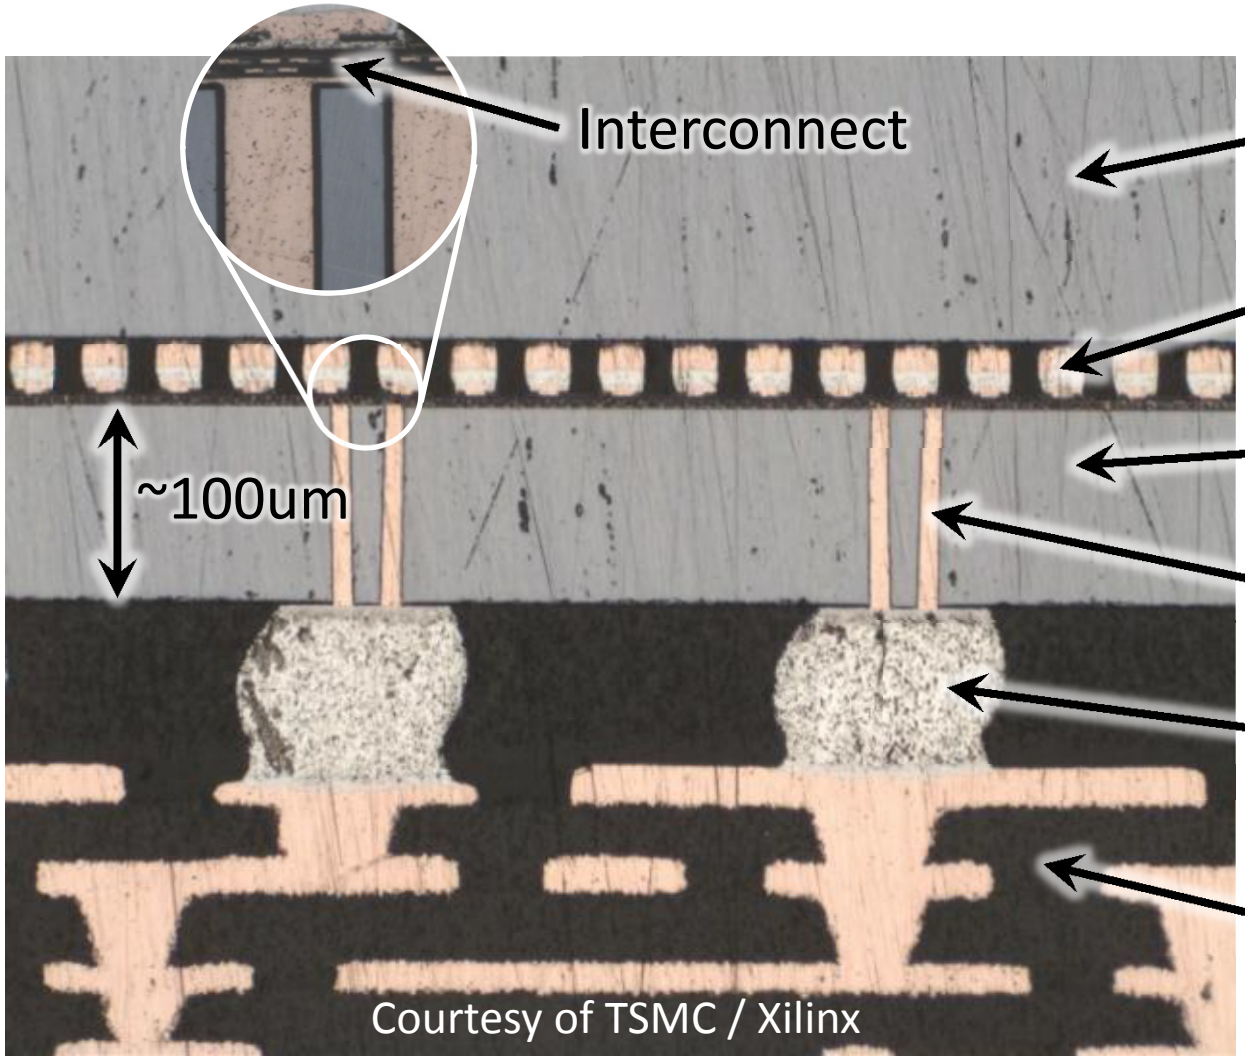

Chip photograph

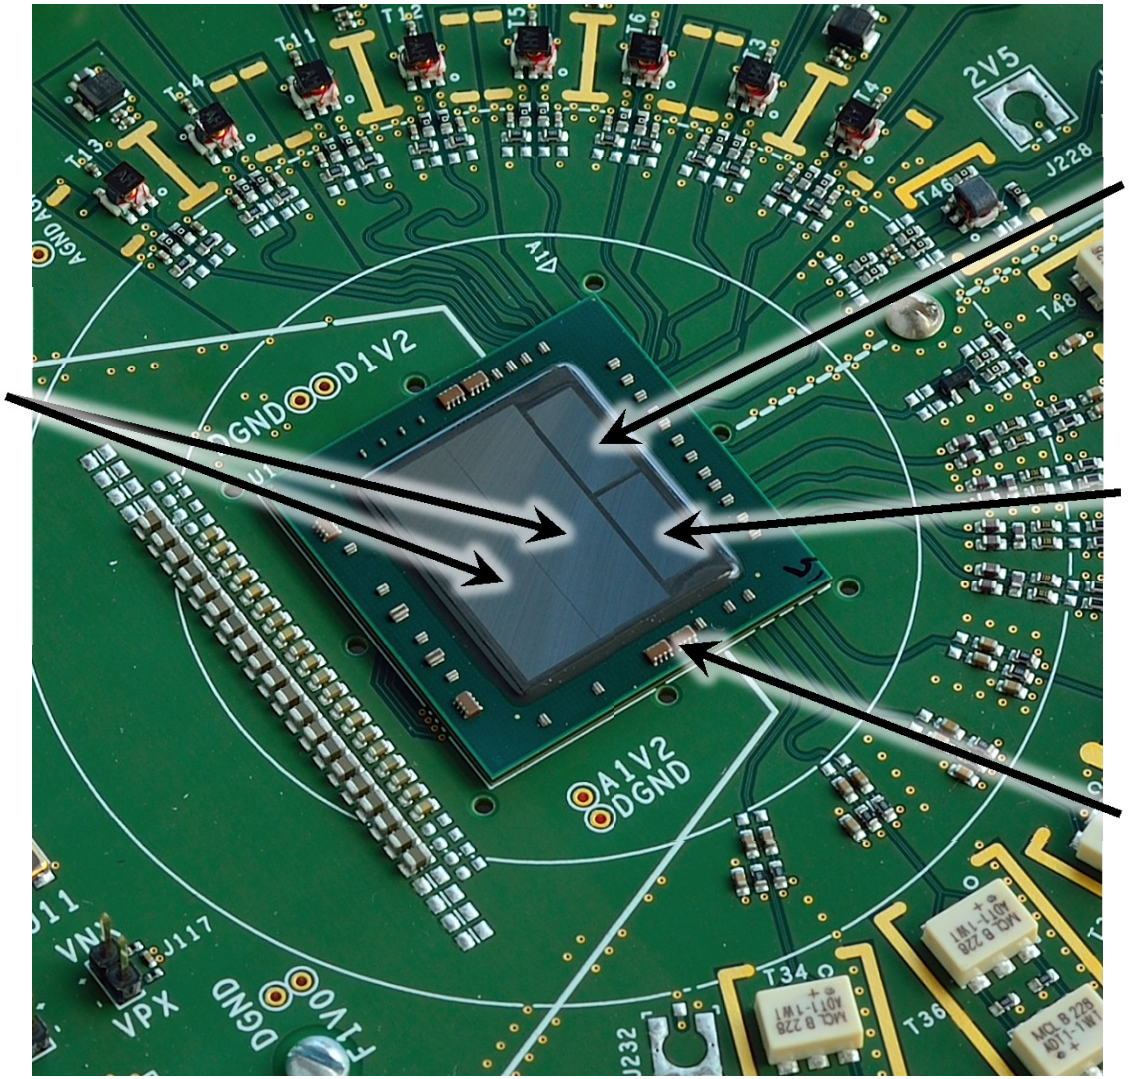

Chip block diagram

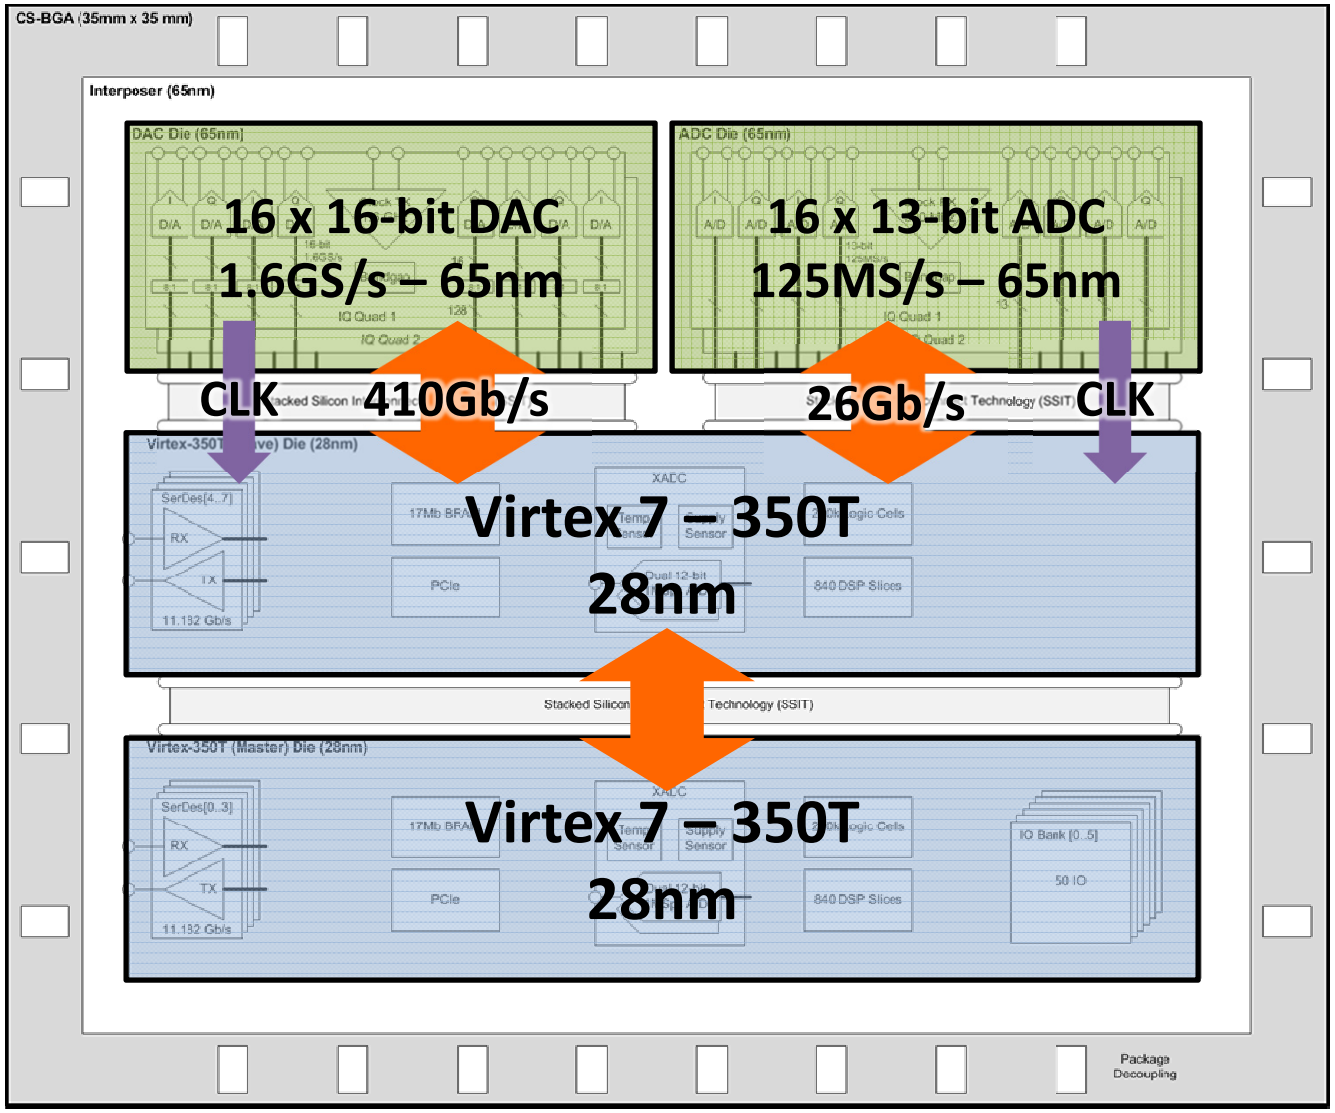

Digital-to-analog isolation strategy

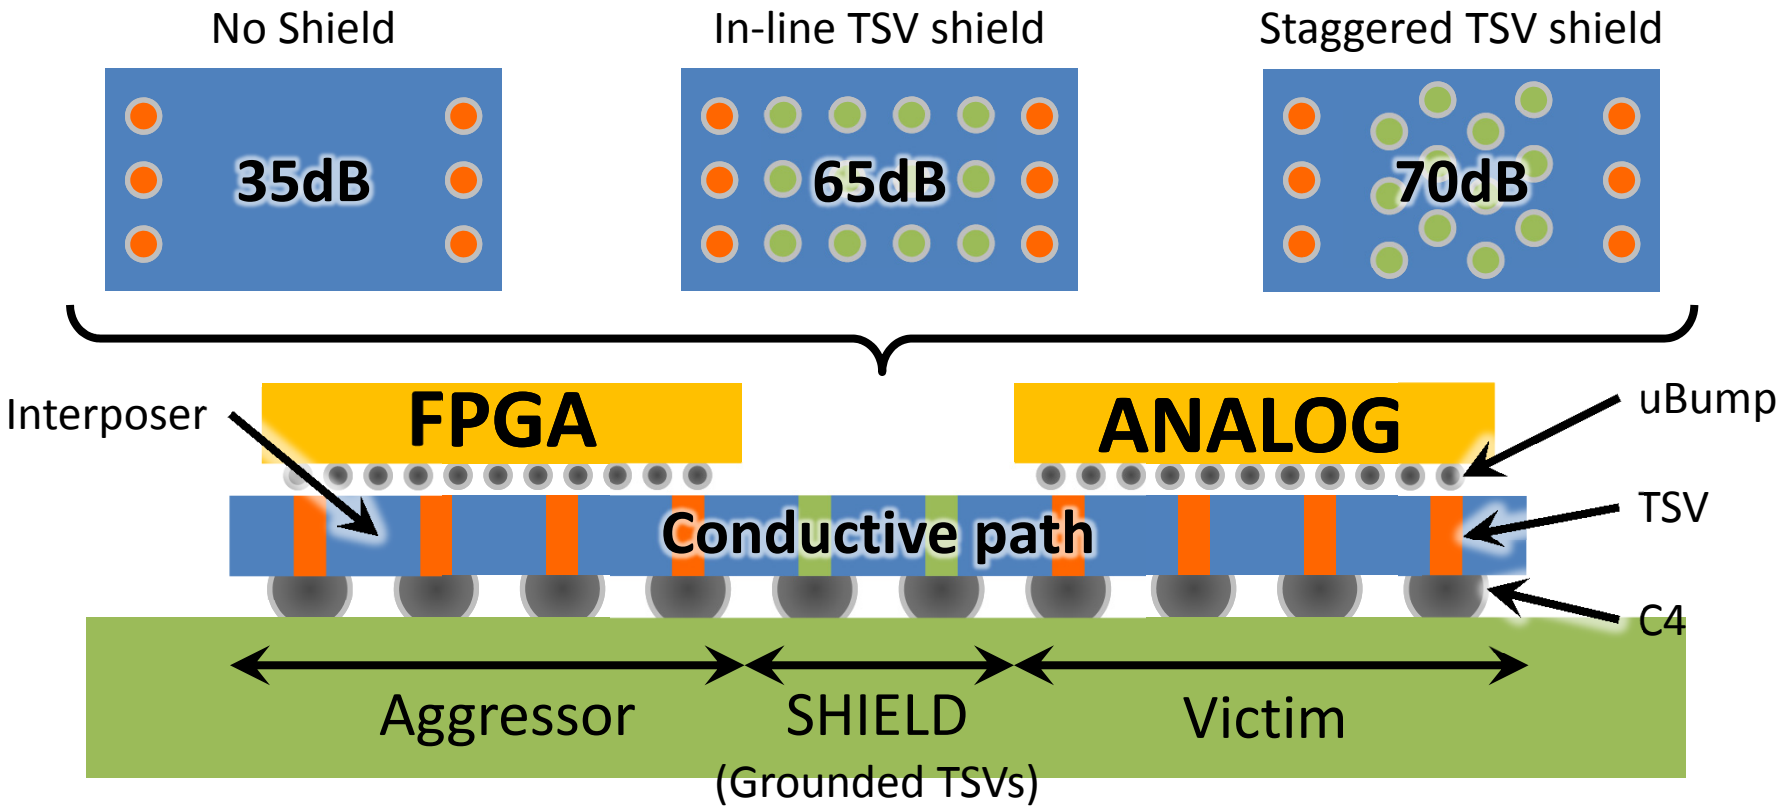

- **Interposer has a resistivity of 20-30 Ω.cm creating coupling**
- **Array of grounded shield TSVs improves interposer isolation from 35dB to 70dB by using staggered scheme**

DAC array block diagram

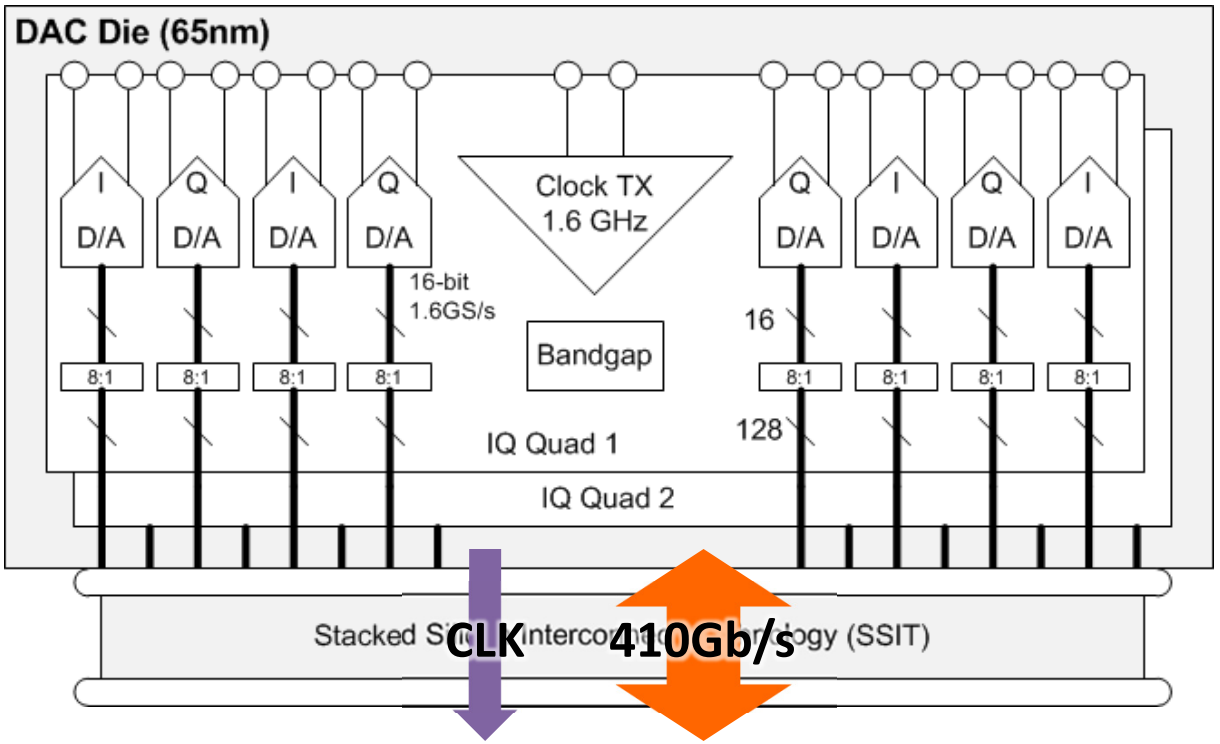

- DAC array is arranged in 2 quads of 4 IQ pair giving a total of 16 DAC instances
- Each quad contains an independent clock receiver and biasing

ADC array block diagram

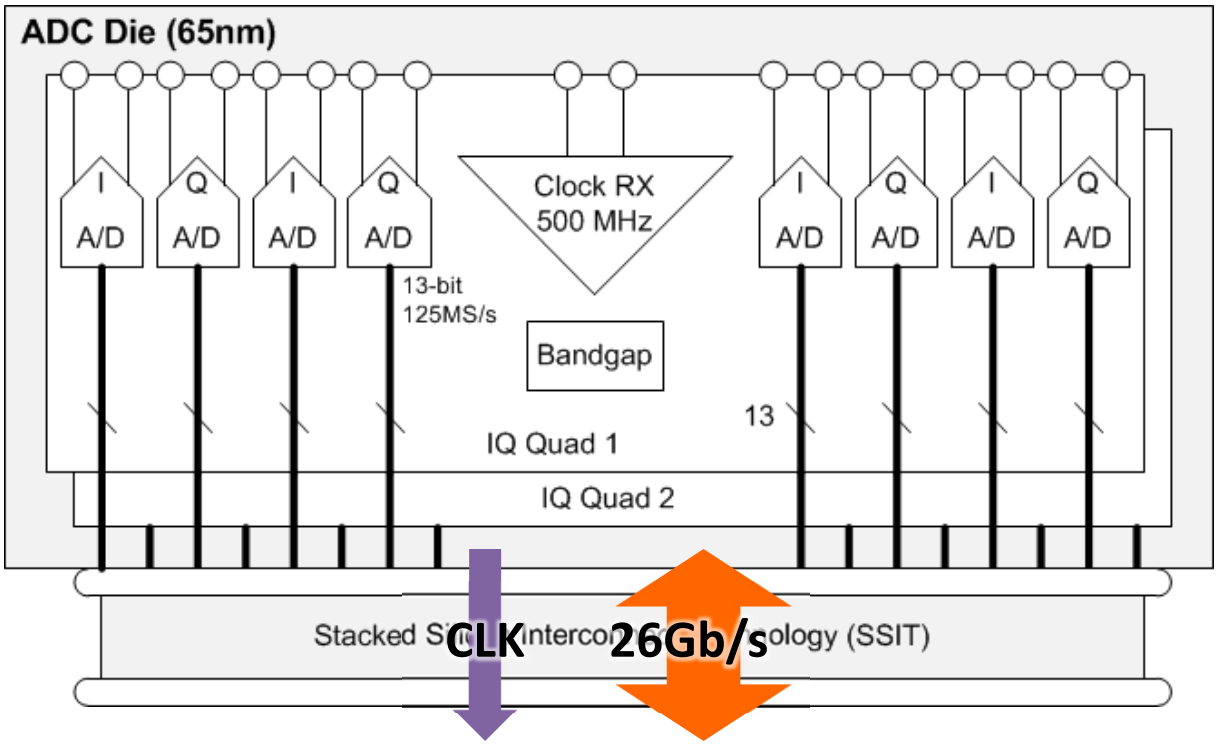

- ADC array arrangement is similar to the DAC offering up to 16 ADC instances
- Channel count can be traded off with sampling speed by using time interleaving

Optimized biasing

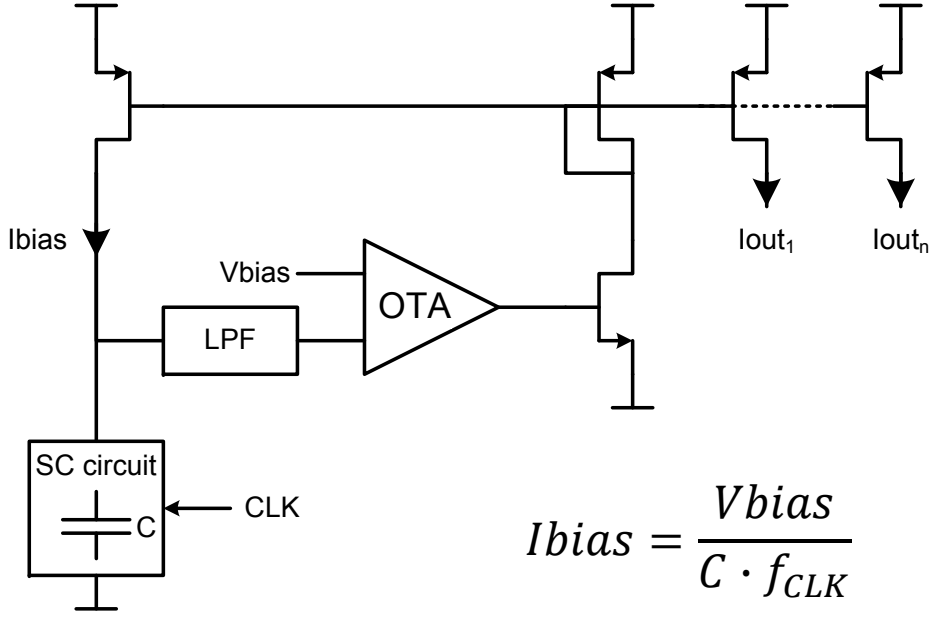

$$Ibias = \frac{Vbias}{C \cdot f_{CLK}}$$

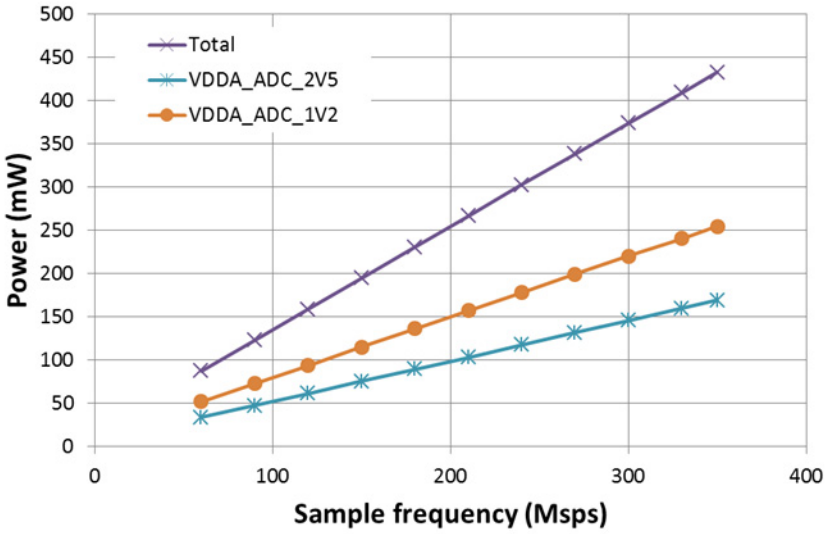

- **Switched capacitor biasing circuit is used to optimize power consumption by tracking frequency AND process variation**

Reconfigurable ADC

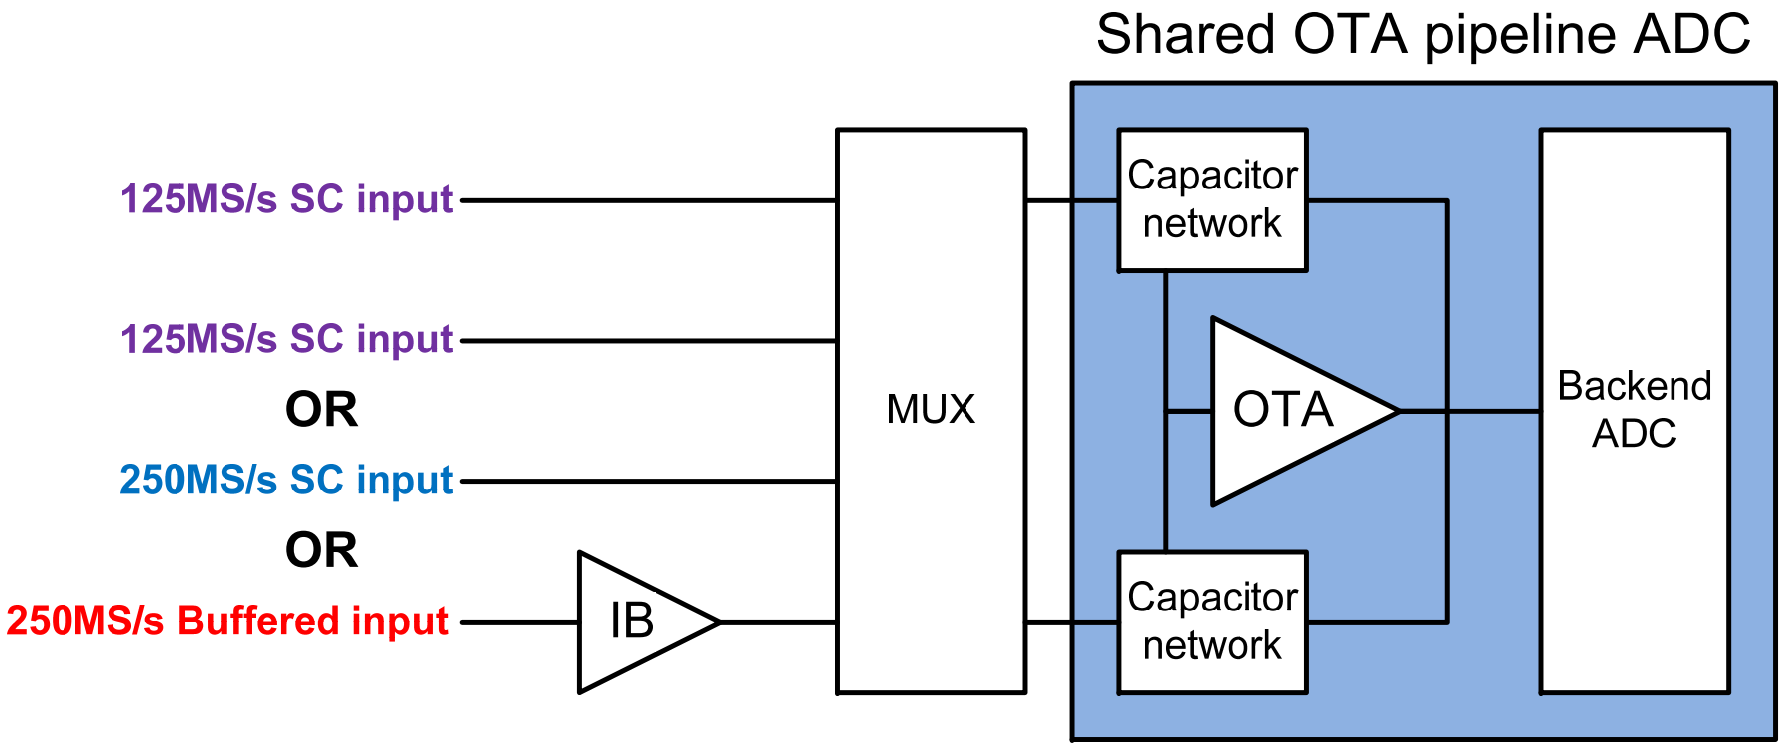

- **ADC is reconfigurable to trade off channel count vs sample rate, including optional input buffer**

DAC measurements

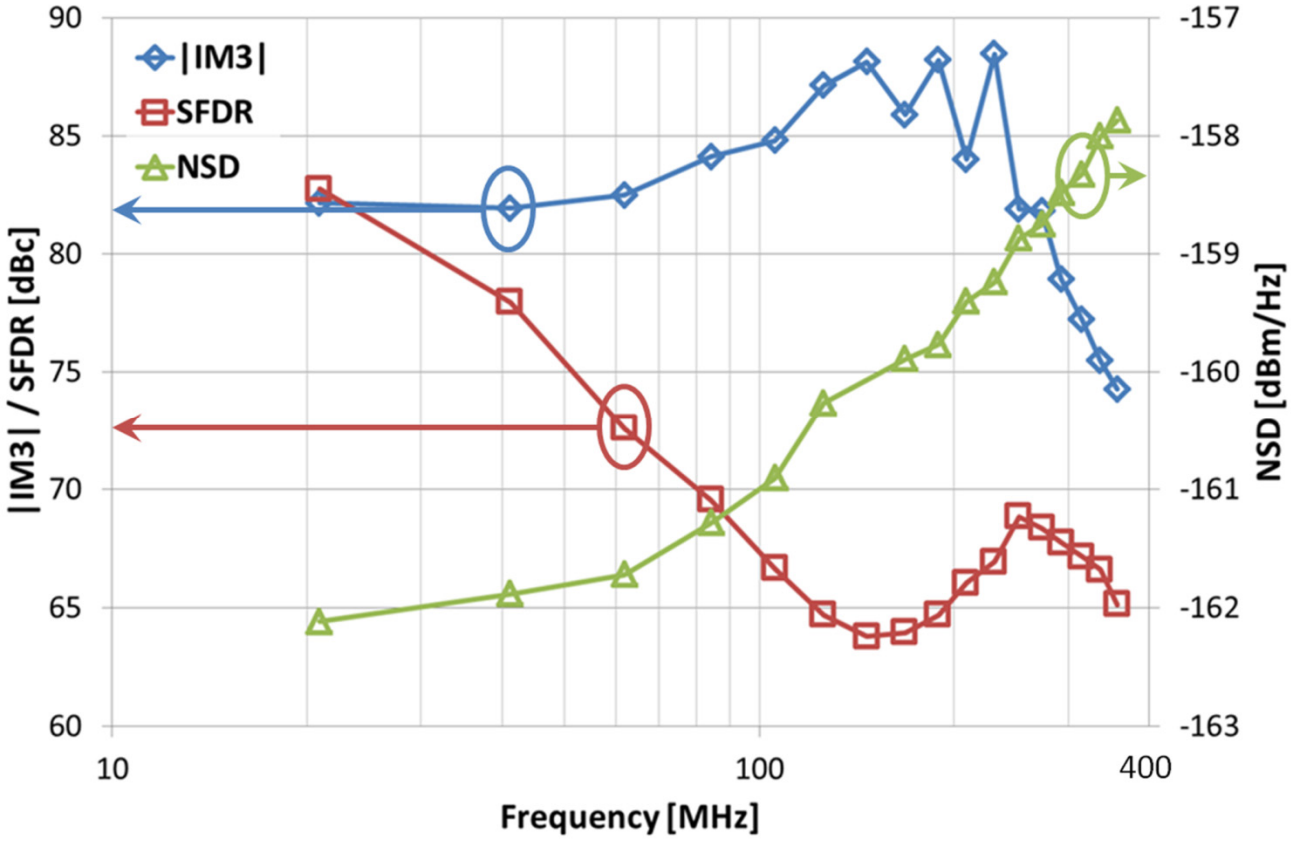

- **Measured DAC dynamic performance (SFDR, IM3 and NSD) versus output signal frequency sampled at 1.6GS/s**

ADC measurements

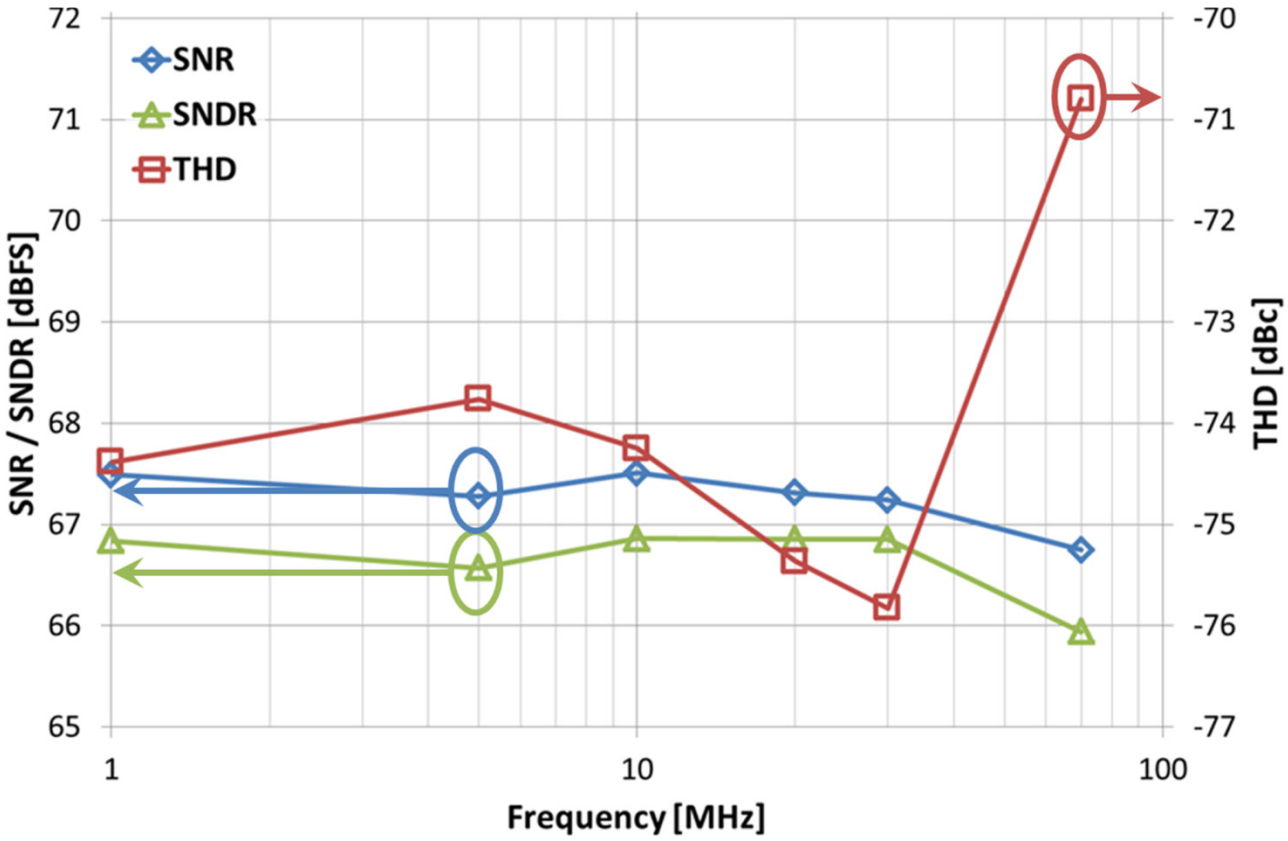

- **Measured ADC dynamic performance (SNR, SNDR and THD) versus input signal frequency sampled at 125MS/s**

4 x interleaved ADC measurements

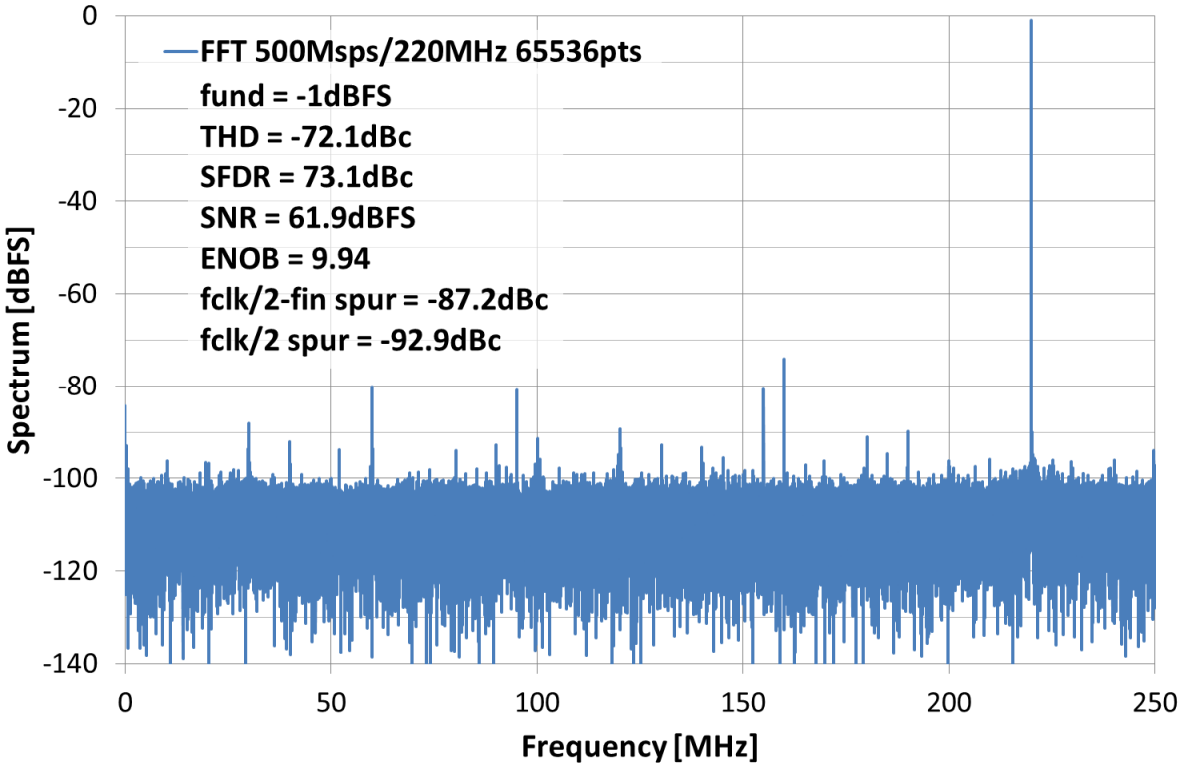

- **Typical FFT spectrum of 4x interleaved ADC for 220MHz input signal frequency sampled at 500MS/s with background calibration of channel mismatch performed on FPGA**

FPGA-to-analog crosstalk measurements

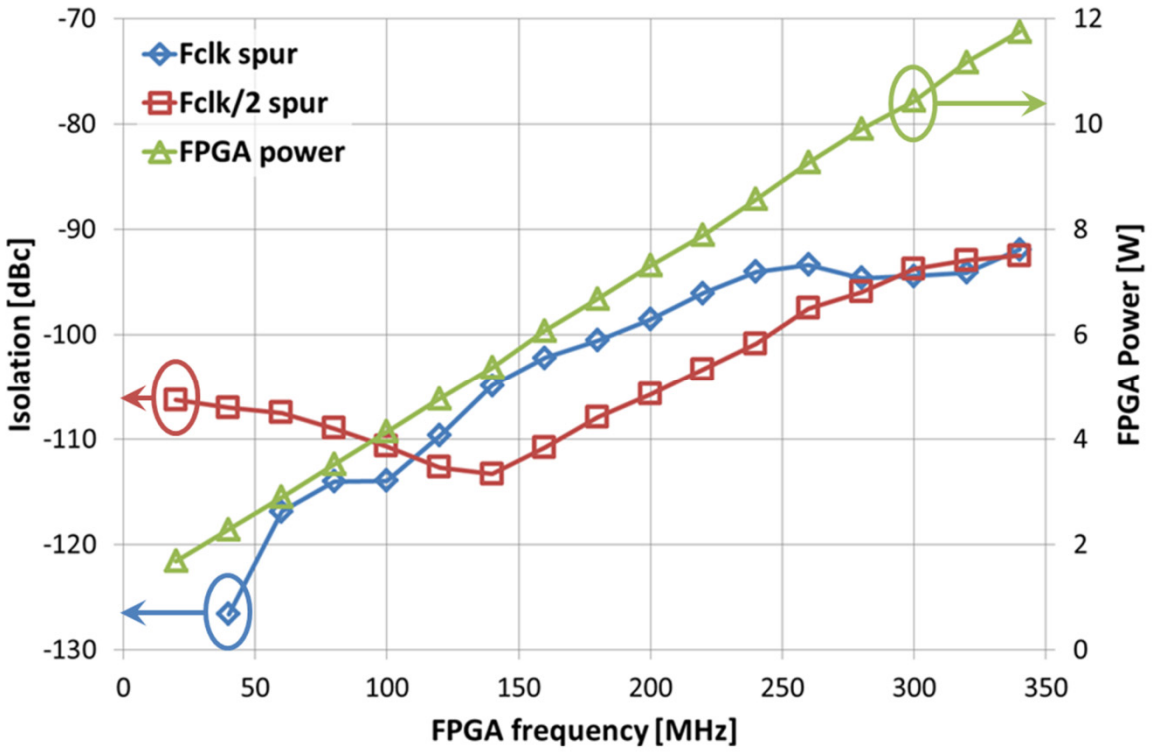

- **Measured FPGA-to-analog isolation while the DAC is synthesizing a 70MHz full scale output tone. 100k DFF within the 28nm logic die and 2k SLL interconnects were simultaneously toggling at the FPGA clock rate**

Summary

DAC Parameter		ADC Parameter		FPGA Resource	
Resolution [bits]	16	Resolution [bits]	13	Logic cells	580480
INL referred to 16b	4	Fclk [MS/s]	125	CLB slices	90700
DNL referred to 16b	4	SNDR [dBFS]	66.6 (64.4)	BRAM [kb]	33840
Fclk [GS/s]	1.6	SNR [dBFS]	67.3 (65.2)	DSP slices	1680
SFDR [dBc] @ 350 MHz	65.1 (59.9)	THD [dBc]	-73.8 (-68.8)	11 Gb/s SerDes	8
IM3 [dBc] @ 300 MHz	-74.3 (-70.2)	SFDR [dBc]	76.6 (69)	IO pins	300
NSD [dBm/Hz]	-157.9 (-157.5)	Unit Power [mW]	134	Interface power [mW/Gb/s]	0.3
Unit Power [mW]	399	FOM [pJ/conv]	0.61	FPGA to analog isolation [dB]	>92

Note: Numbers into brackets indicate minimum value over process, voltage and temperature

- **16 x 16b DAC @1.6GS/s in 65nm**
- **16 x 13b ADC @125MS/s in 65nm (or 8@250MS/s or 4@500MS/s) with optional input buffer**
- **2 x 350T FPGA slices in 28HPL**

Conclusion

- A heterogeneous 3D-IC combining FPGA and high performance data converters has been demonstrated
- Integration of powerful digital processing with highly sensitive analog function was enabled by SSI technology
- Significant power was saved at the data converter interface while outstanding digital-to-analog crosstalk was maintained
- Analog (and digital) reconfigurability was included to address multiple applications